

THE PANCHROMATIC HUBBLE ANDROMEDA TREASURY IV. A PROBABILISTIC APPROACH TO INFERRING THE HIGH MASS STELLAR INITIAL MASS FUNCTION AND OTHER POWER-LAW FUNCTIONS

DANIEL R. WEISZ¹, MORGAN FOUESNEAU¹, DAVID W. HOGG^{2,3}, HANS-WALTER RIX³, ANDREW E. DOLPHIN⁴, JULIANNE J. DALCANTON¹, DANIEL T. FOREMAN-MACKEY², DUSTIN LANG⁵, L. CLIFTON JOHNSON¹, LORI C. BEERMAN¹, ERIC F. BELL⁶, KARL D. GORDON⁷, DIMITRIOS GOULIERMIS^{3,8}, JASON S. KALIRAI⁷, EVAN D. SKILLMAN⁹, BENJAMIN F. WILLIAMS¹

Accepted November 28, 2012

ABSTRACT

We present a probabilistic approach for inferring the parameters of the present day power-law stellar mass function (MF) of a resolved young star cluster. This technique (a) fully exploits the information content of a given dataset; (b) can account for observational uncertainties in a straightforward way; (c) assigns meaningful uncertainties to the inferred parameters; (d) avoids the pitfalls associated with binning data; and (e) can be applied to virtually any resolved young cluster, laying the groundwork for a systematic study of the high mass stellar MF ($M \gtrsim 1 M_{\odot}$). Using simulated clusters and Markov chain Monte Carlo sampling of the probability distribution functions, we show that estimates of the MF slope, α , are unbiased and that the uncertainty, $\Delta\alpha$, depends primarily on the number of observed stars and on the range of stellar masses they span, assuming that the uncertainties on individual masses and the completeness are both well-characterized. Using idealized mock data, we compute the theoretical precision, i.e., lower limits, on α , and provide an analytic approximation for $\Delta\alpha$ as a function of the observed number of stars and mass range. Comparison with literature studies shows that $\sim 3/4$ of quoted uncertainties are smaller than the theoretical lower limit. By correcting these uncertainties to the theoretical lower limits, we find the literature studies yield $\langle\alpha\rangle = 2.46$, with a $1\text{-}\sigma$ dispersion of 0.35 dex. We verify that it is impossible for a power-law MF to obtain meaningful constraints on the upper mass limit of the IMF, beyond the lower bound of the most massive star actually observed. We show that avoiding substantial biases in the MF slope requires: (1) including the MF as a prior when deriving individual stellar mass estimates; (2) modeling the uncertainties in the individual stellar masses; and (3) fully characterizing and then explicitly modeling the completeness for stars of a given mass. The precision on MF slope recovery in this paper are lower limits, as we do not explicitly consider all possible sources of uncertainty, including dynamical effects (e.g., mass segregation), unresolved binaries, and non-coeval populations. We briefly discuss how each of these effects can be incorporated into extensions of the present framework. Finally, we emphasize that the technique and lessons learned are applicable to more general problems involving power-law fitting.

Keywords: stars: luminosity function, mass function – galaxies: star clusters: general methods: statistical

1. INTRODUCTION

The high mass end of the stellar initial mass function (IMF) underpins much of extragalactic astrophysics.

Based on observations made with the NASA/ESA Hubble Space Telescope, obtained from the Data Archive at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.

¹ Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA; dweis@astro.washington.edu

² Center for Cosmology and Particle Physics, New York, University, 4 Washington Place, New York, NY 10003, USA

³ Max Planck Institute for Astronomy, Koenigstuhl 17, 69117 Heidelberg, Germany

⁴ Raytheon Company, 1151 East Hermans Road, Tucson, AZ 85756, USA

⁵ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁶ Department of Astronomy, University of Michigan, 500 Church St., Ann Arbor, MI 48109, USA

⁷ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD, 21218, USA

⁸ Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany

⁹ Minnesota Institute for Astrophysics, University of Minnesota, 116 Church Street SE, Minneapolis, MN 55455, USA

Stars more massive than a few solar masses are largely responsible for most chemical enrichment, dominate the spectral energy distribution blueward of $\sim 1\mu\text{m}$ for all star-forming galaxies, and are presumed to dominate stellar feedback processes on galactic scales. Consequently, the exact numbers and mass distributions of high mass stars are central to the interpretation of integrated light from distant galaxies, chemical evolution models, the frequency of core-collapse supernovae, the evolution of star formation rates over cosmic time, and the efficiency of star formation on galactic scales (e.g., Leitherer et al. 1999; Bruzual & Charlot 2003; Tinsley 1980; Smartt 2009; Madau et al. 1996; Schmidt 1959; Kennicutt 1989; Leroy et al. 2008). Further, the form of the IMF and its potential sensitivity to local environment carries significant implications for how stars form from dense molecular cores and how the individual stars affect the properties of stellar clusters (e.g., Elmegreen 1999, 2004; Weidner & Kroupa 2005; Zinnecker & Yorke 2007; McKee & Ostriker 2007; Weidner et al. 2010; Portegies Zwart et al. 2010; Bastian et al. 2010).

The IMF can be sensibly parameterized in a number of ways (e.g., Salpeter 1955; Miller & Scalo 1979; Scalo

1986; Chabrier 2003), in particular as a set of power-laws. Following Kroupa (2001), we adopt the following parameterization of the IMF:

$$\Phi(M) \equiv \left\langle \frac{dN}{dM} \right\rangle = c_i M^{-\alpha_i}, \quad M_{a,i} \leq M \leq M_{b,i}, \quad (1)$$

where c_i are chosen to ensure continuity, $\Phi(M)$ is normalized such that

$$\int \Phi(M) dM = 1 M_{\odot}, \quad (2)$$

and $\vec{\alpha} \equiv \alpha_i, i = 1, 3$ is slope of each power-law component within specified mass interval such that

$$\begin{aligned} \alpha_1 &= 0.3, & 0.01 \leq M/M_{\odot} < 0.08 \\ \alpha_2 &= 1.3, & 0.08 \leq M/M_{\odot} < 0.5 \\ \alpha_3 &= 2.3, & 0.5 \leq M/M_{\odot} \leq M_{\max}, \end{aligned}$$

with M_{\max} being the maximal mass of a star that could have been formed. The value of M_{\max} can either be set by the maximal mass at which the stellar lifetime exceeds the (finite) age of the cluster, or by the star- or cluster-formation process itself (e.g., Elmegreen 2002, 2004; Weidner & Kroupa 2005; Goodwin & Kouwenhoven 2009; Myers et al. 2011; Krumholz et al. 2010, 2011). The high mass portion of the IMF (i.e., α_3) sets essentially all observables (e.g., luminosity, color, chemical enrichment, etc.) in extragalactic contexts, making analysis of a single power-law a reasonable approximation in most environments outside the Galaxy. We will therefore use shorthand notation $\alpha \equiv \alpha_3$ throughout this paper.

Despite its widespread importance, the IMF for stars above $\sim 1 M_{\odot}$ remains insufficiently constrained by observations. As shown in Figure 1 and listed in Table 1, measurements of the IMF slope from studies of resolved star clusters exhibit a scatter of ~ 0.5 dex. Additionally, various compilation ‘ α -plots’ from the literature derive mean IMF slopes with variations of ~ 0.4 dex between different compilations (e.g., Scalo 1998; De Marchi et al. 2010; Bastian et al. 2010). Taken at face value, the observed scatter in individual IMF measurements and variations in the average values of different compilations introduce significant systematics into interpreting the higher redshift universe (e.g., Cool et al. 2008; Conroy et al. 2009; Narayanan & Davé 2012), and further inhibit differentiation among star formation models that result in a universal IMF (e.g., McKee & Ostriker 2007; Myers et al. 2011) and those that are sensitive to environmental conditions (e.g., Weidner & Kroupa 2005), each of which present a drastically different view of how galaxies evolve (e.g., Pflamm-Altenburg & Kroupa 2009; Haas & Anders 2010).

However, at present, it remains unclear whether the precision and accuracy of IMF slope measurements should be taken at face value, which either implies a formal rejection of a universal high mass IMF slope or whether it should be interpreted as a reflection of stochasticity in the measurements.

Currently, the majority of high mass IMF studies directly count stars or measure luminosity functions in single or small sets of clusters, with each employing different




Figure 1. Selected literature measurements for the high mass ($M \gtrsim 1 M_{\odot}$) IMF slope from resolved star counts (see Table 1). Points have been placed on the x-axis at the mean mass of the range covered by a given study. The thin grey lines indicate the span of the full mass range. Error bars in the y-direction reflect the quoted 1- σ uncertainties. The heterogeneity in the MF recovery techniques makes it difficult to discern whether the ~ 0.5 dex of scatter of the data is due to physical or observational challenges and further convolutes the significance of the uncertainties.

observational and IMF recovery techniques. Many studies selectively address critical issues including corrections for binary stars, spatially dependent observational completeness, and uncertainties in individual stellar masses. Consequently, whether the broad range of measurements in Figure 1 are due to intrinsic physical effects or are the result of systematics remains an open and important question (see discussions in Scalo 1986; Elmegreen 1999; Kroupa 2002; Scalo 2005; Elmegreen 2009; Bastian et al. 2010).

In practice, we can only measure the present day mass function (MF) of evolved clusters, in which the observable stellar masses are limited at the high mass end by the turnoff mass corresponding to the age of the cluster, although the object of ultimate importance is the entire IMF. For a simple stellar population, i.e., a star cluster, the measured MF slope is identical to the IMF slope (e.g., Elmegreen & Scalo 2006) modulo certain dynamical (e.g., mass segregation) and observational (e.g., completeness) effects. Consequently, our understanding of the IMF critically depends on our ability to accurately and precisely constrain the MF.

Systematically counting the relative numbers of individual stars (with $M \gtrsim 1 M_{\odot}$) in a large collection of young clusters has long been advocated as the optimal approach to constraining the high mass MF (e.g., Scalo 1986; Kroupa 2002; Scalo 2005; Bastian et al. 2010; Kroupa et al. 2011). As improved datasets of resolved stars have emerged (e.g., Zaritsky et al. 1997; Massey et al. 2006; Holtzman et al. 2006; Dalcanton et al. 2009, 2012; Bianchi et al. 2012), it is crucial that the analysis tools for constraining the high mass MF are in place to fully exploit the information content of the data and to properly account for all known uncertainties.

Broadly speaking, the practical procedure of obtaining a MF constraint for a resolved young cluster consists of three steps. First, one obtains individual mass estimates for the stars, either based on photometry in conjunction with isochrones, or from spectroscopy. Second, one de-

termines the completeness of the sample for which mass estimates were obtained. Finally, one fits a MF model (e.g., Equation 1) to the data to obtain confidence limits on the slope parameter, α . In its usual implementation, this approach uses χ^2 minimization of a power-law model fit to a binned or cumulative representation of the mass function, $N(M)$. The fit is taken over the masses determined to be sufficiently complete and the error bars on each mass bin reflect Poisson statistics of the number of stars in the bin.

However, there are severe limitations in this traditional approach. Correlations between the number of stars per bin and the adopted bin size can introduce significant biases into the resulting MF slope and associated uncertainties (e.g., Maíz Apellániz & Úbeda 2005; Cara & Lister 2008). Further, such an approach does not provide an objective framework for folding in observational effects such as completeness and mass uncertainties, which are critical to accurate MF constraints. More robust approaches such as Monte-Carlo simulations, color-magnitude diagram fitting, non-parametric kernel estimators, and maximum likelihood techniques (e.g., Elson et al. 1989; Gennaro et al. 2011; Dolphin 2002; Vio et al. 1994; Tarrab 1982) alleviate some of these issues. However, their reliability has not been shown in large, diverse sets of clusters, the interpretation of the error bars is often ambiguous, and there remain statistical short-comings in even the most advanced conventional approaches (e.g., noise amplification in maximum likelihood estimations; Gould et al. 1997). In short, traditional approaches to MF reconstruction are inadequate for a systematic investigation of the high mass stellar IMF.

The main goal of this paper is to lay out a forward modeling technique that remedies at least some of the shortcomings of the traditional approach. It should result in constraints on the high mass MF that come close to exploiting the full information content of the data and that explicitly account for known observational uncertainties. In particular, by laying out a forward modeling formalism, we aim to overcome the following limitations of many previous analyses by: (1) accounting for the mass estimate errors of individual stars that may arise from the conversion of observed fluxes into stellar masses; (2) avoiding binning the measurement in mass, as the most massive bins are unavoidably sparsely populated; (3) dealing explicitly with completeness functions that may not be a simple step function of stellar mass, and instead be either gradual or varying as a function of position; (4) implementing a rigorous way to derive the joint posterior distribution function of the MF parameters (e.g., α and M_{\max}) in light of evidence for the data and its uncertainties and possibly other prior information.

In this paper, we present the probabilistic framework for inferring the parameters of the high mass MF of a young stellar cluster. We carefully walk through the necessary mathematics and interlace illustrative examples using simulated clusters. For clarity in the derivations and examples, we have made several simplifying assumptions. In particular, dynamical effects such as mass segregation, ejection, and relaxation are not included in the present paper, although their role in interpreting the observed MF as the IMF is indisputably important.

The methodology presented in this paper is motivated in part by the Panchromatic Hubble Andromeda Treasury program (PHAT; Dalcanton et al. 2012). This Hubble Space Telescope multi-cycle treasury program is mapping $\sim 1/4$ of M31’s star forming disk, providing near-UV through near-IR imaging of $\sim 10^8$ resolved stars. Based on analysis of the Year 1 data, we anticipate the survey will resolve $\gtrsim 1000$ young clusters (i.e., ages $\lesssim 100$ Myr; or a main sequence upper mass limit of $\gtrsim 5 M_{\odot}$; Johnson et al. 2012) with sensitivity down to stars with $M \sim 1\text{--}3 M_{\odot}$, which will enable a large scale study of the high mass stellar MF. Although the terminology ‘high mass’ is often reserved for stars with $M \gtrsim 10\text{--}15 M_{\odot}$, common IMF parameterizations (cf. Equation 1) suggest that all stars above $\sim 1 M_{\odot}$ are drawn from the same underlying IMF, indicating that conventionally defined intermediate and massive stars can be used to learn about their common IMF. Throughout the paper, we use these numbers as guides for our simulations, but emphasize that the methodology developed in this paper can readily be generalized to measure the MF of any resolved cluster and for any parameterized MF model.

The remainder of this paper is organized as follows. In §2, we lay out the general probabilistic framework for inferring the MF of a young stellar cluster. We then apply this technique to highly idealized mock data (e.g., no mass uncertainties, perfectly known completeness) and present illustrative results in §3. In §3.2 and §3.3, we derive the theoretical lower limits for recovered MF slope precision and compare the results to select literature studies. We then consider the more complex case of masses with finite mass uncertainties in §4, and explore the case where the stellar mass distributions are log-normal functions in §4.1. We then consider the case of linear completeness functions in §5. Finally, we discuss several caveats to the MF models adopted in this paper (e.g., the influence of cluster dynamics, non-coeval populations, unresolved binaries) in §6.

2. PROBABILISTIC FRAMEWORK FOR MODELING THE STELLAR MASS FUNCTION OF A CLUSTER

We begin by illustrating the scope of the problem at hand. Consider the IMF of a hypothetical cluster as shown in Figure 2. We presume that the high mass IMF regime of this cluster consists of some number of stars, whose mass distribution follows a power-law with a slope of α , and whose masses are bounded by M_{\min} and M_{\max} , which is the mass of the most massive star that could have formed in the cluster, and may not be the same for different clusters. This ‘zero-age’ state of the cluster is the desired dataset to make direct inferences about its stellar IMF.

However, there are several considerations that only allow us to observe the present day MF of this cluster. First, due to stellar evolution, the most massive star(s) have likely disappeared, and only stars with masses up to $M_{\max, \text{obs}}$ can be observed at the present day. Second, observational completeness (due to stellar crowding, for example) only permits observations of stars more massive than some lower mass limit, M_{comp} . Third, the translation of observed flux into stellar mass provides only an estimate of a star’s true mass. That is, the ‘observed’ mass of a star could be an over- or underestimate of a star’s true mass. Further, the observed mass of a star




Figure 2. A schematic representation of the Kroupa IMF over the entire stellar mass spectrum for a hypothetical star cluster (Equation 1). Here, N_* (red shading) represents the total number of stars in the cluster and N_{pred} is the number of stars expected in the observed mass range bounded by M_{comp} and $M_{\text{max,obs}}$ (blue shaded region), i.e., the MF. The upper end of a hypothetical power-law IMF is defined to be between M_{min} and M_{max} , which is the maximal mass of a star that could have formed in the cluster. M_{lim} is the upper bound of possible observed stellar masses, as set by stellar physics. The aim of this paper is to outline an approach for inferring the values of α , M_{max} , and N_{pred} of a young cluster, given a set of N observed stellar masses.

is bounded by M_{comp} and M_{lim} , the maximum mass allowed by stellar evolution models, which may or may not reflect the upper stellar mass limit in nature. Each of these effects must be carefully considered in order to make an accurate measurement of the present day MF of a cluster, a necessary step to make any meaningful statements about a cluster's stellar IMF.

For clarity in exposition, we will begin with the simplest case where the data, \vec{d} , consists of a set of N stars with perfectly known stellar masses, $\{M_i\}$, and the completeness function is also perfectly known. Given the data, \vec{d} , we aim to place constraints on the slope of the MF, α , the maximum mass of a star that could have formed, M_{max} , and the expected number of observed stars, N_{pred} , which normalizes the amplitude of the MF for this cluster (see Figure 2). Throughout this paper, we assume that we have the observed completeness function (either as a function of mass or luminosity) in hand. In practice, the completeness is a function of both the stellar fluxes (i.e., through the flux limit of the observations) and the local image crowding, and is determined through extensive artificial star tests. Given that most cluster observations, certainly those beyond the Milky Way, are crowding limited, it will always be more challenging to detect and accurately catalog a particular star if it is located near the center of a cluster rather than in the outskirts. For the purposes of this paper, we make the simplifying assumption that the completeness function is independent of position.

2.1. Deriving the General MF Posterior Probability Distribution Function

We now lay out the general probabilistic framework for constraining the MF of a young cluster. We start by describing the MF of the cluster (cf. Equation 1) as:

$$\Phi(M) = N_* p_{\text{MF}}(M | \vec{\theta}), \quad (3)$$

where N_* is the total number of stars in the cluster, $p_{\text{MF}}(M | \vec{\theta})$ is the probability distribution function (PDF) for the mass of any star in the cluster, and M is the true mass of the star; the MF is described by the parameters $\vec{\theta} = \{\alpha, M_{\text{max}}\}$, where α and M_{max} refer to the high mass portion of the MF, i.e., a single sloped power law such as described by α_3 in Equation 1.

The PDF for the mass of an observed star, i.e., those that are massive (luminous) enough to end up in our dataset, is given by

$$p_{\text{MF}_o}(M | \vec{\theta}, \text{obs}) \equiv \frac{p_{\text{MF}}(M | \vec{\theta}) p(\text{obs} | M)}{p(\text{obs} | \vec{\theta})}, \quad (4)$$

where $p(\text{obs} | M)$ is the probability of observing a star given its true mass M , i.e., the completeness function, and $p(\text{obs} | \vec{\theta})$, the normalization necessary to make $p_{\text{MF}_o}(M | \vec{\theta}, \text{obs})$ a PDF. This normalization can be written as

$$p(\text{obs} | \vec{\theta}) = \int p_{\text{MF}}(M | \vec{\theta}) p(\text{obs} | M) dM. \quad (5)$$

In this model, the expected number of stars in the data list, N_{pred} , sets the amplitude normalization of the MF, and can be expressed as

$$N_{\text{pred}} = N_* \int_0^\infty p_{\text{MF}_o}(M | \vec{\theta}, \text{obs}) dM. \quad (6)$$

However, given that we do not know N_* *a priori* and that much of the MF is, as in many practical applications, well below our detection limit, we cannot explicitly solve for N_{pred} , and therefore we will treat N_{pred} as a model parameter that is independent of $\vec{\theta}$.

So far, this model makes predictions for the probability distribution for the mass of any one observed star, $p_{\text{MF}_o}(M | \vec{\theta}, \text{obs})$, for a given set of parameters, $\vec{\theta}$, and for the expected number of observed stars, N_{pred} . Assuming that each of the N observed stars in the cluster is drawn identically and independently from the underlying mass distribution, we can write the probability of measuring the set of N masses, $\{M_i\}$, as:

$$p_{\text{M}}(\{M_i\} | \vec{\theta}, \{\text{obs}_i\}) = \prod_{i=1}^N p_{\text{MF}_o}(M_i | \vec{\theta}, \text{obs}_i). \quad (7)$$

At the same time, we must also include the probability of actually observing N stars in the cluster, as N itself is a datum. This is the Poisson probability of observing N stars given an expectation of N_{pred} stars and can be written as

$$p_{\text{Poisson}}(N | N_{\text{pred}}) = \frac{N_{\text{pred}}^N e^{-N_{\text{pred}}}}{N!}. \quad (8)$$

We now have the two expressions that give the probability of the data ($\{M_i\}, N$) given model parameters ($\vec{\theta}, N_{\text{pred}}$) and the fact that data were observed, i.e., ‘obs’. However, we wish to infer the inverse, namely the values of $\vec{\theta}$ and N_{pred} given a set of N stellar mass measurements. Using Bayes’s theorem, we can then write down the probabilities of the model parameters given the data, i.e., the posterior PDFs (pPDFs) for $\vec{\theta}$ and N_{pred} as:

$$\begin{aligned}
p_{\text{post}}(\vec{\theta} | \{M_i\}, \text{obs}) &= p_{\text{M}}(\{M_i\} | \vec{\theta}, \{\text{obs}_i\}) p_{\text{prior}}(\vec{\theta}) \\
&= \prod_{i=1}^N p_{\text{MF}_o}(M_i | \vec{\theta}, \text{obs}_i) p_{\text{prior}}(\vec{\theta}), \quad (9)
\end{aligned}$$

and

$$\begin{aligned}
p_{\text{post}}(N_{\text{pred}} | N) &= p_{\text{poisson}}(N | N_{\text{pred}}) p_{\text{prior}}(N_{\text{pred}}) \\
&= \frac{N_{\text{pred}}^N e^{-N_{\text{pred}}}}{N!} p_{\text{prior}}(N_{\text{pred}}), \quad (10)
\end{aligned}$$

where $p_{\text{prior}}(\vec{\theta})$ and $p_{\text{prior}}(N_{\text{pred}})$ reflect any prior or extraneous information or constraints on $\vec{\theta}$ and N_{pred} , respectively. In Equations 9 and 10 we have presumed that the ‘evidence’ terms, $p_{\text{prior}}(\{M_i\})$ and $p_{\text{prior}}(N)$, are both constant, and have omitted them. In practice, the independence of these two pPDFs allows us to compute them separately, and then simply multiply them together to get the general pPDF for the MF:

$$\begin{aligned}
p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \{M_i\}, N, \text{obs}) &= \\
p_{\text{post}}(\vec{\theta} | \{M_i\}, \text{obs}) p_{\text{post}}(N_{\text{pred}} | N). \quad (11)
\end{aligned}$$

The above derivation is limited to a single power-law slope, where the stars are well-above any break or turnover in the MF. Of course, in cases where data on stars well below $1 M_{\odot}$ are available, the same formalism carries through, just replacing α with $\tilde{\alpha}$.

3. DERIVING MASS FUNCTION CONSTRAINTS FROM IDEALIZED DATA

We now show how to explicitly constrain MF parameters given: (1) the pPDF from Equation 11; (2) evidence for the data; (3) and any prior information on $\vec{\theta}$ and N_{pred} . In practice, the most informative prior on $\vec{\theta}$ frequently comes from M_{max} , where astrophysical constraints may exist, especially if there is independent information on the age of the population. Other general, although less physically informative, priors are requirements that the number of stars is positive and that α does not take on unreasonably extreme values that have been ruled out by previous observations, e.g., $\alpha < -5$.

For the purposes of this initial idealized exercise, we adopt a simple ‘boxcar’ form for the probability of observing a star, given its true mass (i.e., the completeness function),

$$p(\text{obs} | M) = \begin{cases} 1, & 0.5 M_{\odot} < M_{\text{comp}} \leq M \leq M_{\text{lim}} \\ 0, & \text{otherwise} \end{cases}, \quad (12)$$

where M_{lim} is the maximum observable mass as set by stellar evolution, M_{comp} is the minimum mass as determined by the observational completeness limit. For simplicity, we will assume that the MF is a single power-law over the interval $M_{\text{comp}} \leq M \leq M_{\text{lim}}$ and that M_{comp} is the same value for all locations in the cluster; we discuss the case of spatially varying completeness in §6. In practice, $p(\text{obs} | M)$ is the completeness function of the data.

So far, we have assumed that the uncertainties on masses are negligible. It follows that a star’s observed mass PDF is a delta function, and therefore it is directly interchangeable with a delta function PDF for the true mass of a star. Consequently, we can use $p_{\text{MF}_o}(M | \vec{\theta}, \text{obs})$ as the generative model for the observations such that

$$\begin{aligned}
p_{\text{MF}_o}(M | \vec{\theta}, \text{obs}) &= p(\text{obs} | M) c_{\text{MF}_o}(\vec{\theta}) M^{-\alpha} = \\
&\begin{cases} c_{\text{MF}_o}(\vec{\theta}) M^{-\alpha}, & M_{\text{comp}} \leq M \leq M_{\text{max}} \\ 0, & \text{otherwise} \end{cases}, \quad (13)
\end{aligned}$$

where the normalization $c_{\text{MF}_o}(\vec{\theta})$ is given by

$$c_{\text{MF}_o}^{-1}(\vec{\theta}) = \int_0^{\infty} p(\text{obs} | M) M^{-\alpha} dM = \int_{M_{\text{comp}}}^{M_{\text{max}}} M^{-\alpha} dM. \quad (14)$$

In this form, the range of the MF model is set by the completeness limit on the low mass end, and the mass of the most massive star to have formed on the upper end.

Using the explicit expressions for each term in Equation 11, we can now re-write the pPDF for this idealized case as:

$$\begin{aligned}
p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \{M_i\}, N, \text{obs}) &= \\
p_{\text{poisson}}(N | N_{\text{pred}}) \sum_{i=1}^N (c_{\text{MF}_o}(\vec{\theta}) M_i^{-\alpha}) p_{\text{prior}}(\vec{\theta}) p_{\text{prior}}(N_{\text{pred}}), \quad (15)
\end{aligned}$$

In practice we calculate the natural logarithm of the pPDF for computational ease. Also, note that in this idealized limit, $p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \{M_i\}, N, \text{obs}) = 0$ in the event that $M_i < M_{\text{comp}}$ or $M_{\text{max}} > M_{\text{lim}}$.

3.1. Sampling the the pPDF with a Markov Chain

To place constraints on the model parameters of interest, we sample the pPDF using a Markov chain Monte Carlo (MCMC) algorithm. MCMC techniques provide an efficient discrete sampling and clear interpretation of multi-dimensional spaces, such that the density of samples is highest around the most probable parameters and lower in regions of less probable values (e.g., Gelman et al. 1996). As a result, MCMC techniques are able to produce well-defined uncertainties for both the one dimensional (marginalized) distributions (e.g., $p(\alpha | \{M_i\}, N, \text{obs})$, $p(M_{\text{max}} | \{M_i\}, N, \text{obs})$) and reveal degeneracies between two dimensional (joint) distributions (e.g., $p(\alpha, N_{\text{pred}} | \{M_i\}, N, \text{obs})$). An MCMC approach also permits efficient sampling in cases where more parameters are needed to adequately model the MF, e.g., such as when using a multi-component power-law to characterize data containing lower mass stars.

In this analysis, we use the MCMC sampler `emcee`¹⁰ as described in Foreman-Mackey et al. (2012). `emcee` is a pure-Python implementation of the affine invariant MCMC ensemble sampler developed by Goodman & Weare (2010). This sampler explores parameter space through a set of ‘walkers’. For each MCMC increment,

¹⁰ <http://danfm.ca/emcee>

each walker takes a step in parameter space by choosing another walker and moving along a line in parameter space that connects itself to the other walker. The size of the step is chosen stochastically, allowing for both interpolation and extrapolation, and the choice in step for each walker is based on the co-variance of the set of walkers. After each step, the pPDF is evaluated for the new set of parameters. Steps that increase the probability are always accepted, whereas steps that result in a lower probability are *sometimes* accepted. For a sufficiently large number of steps, the ensemble of walkers samples parameter space with a frequency proportional to the pPDF. In a rough sense, this approach is akin to having a set of parallel Metropolis-Hastings MCMC chains (Metropolis et al. 1953; Hastings 1970), but is much more efficient than Metropolis-Hastings sampling, as measured by the autocorrelation time (i.e., the time spent per function call; see discussions in Goodman & Weare 2010, Foreman-Mackey et al. 2012, and the Appendix of this paper).

`emcee` requires minimal initial input from the user. One must select a number of walkers and their initial conditions, and designate the number of burn-in steps and length of each chain. Following the recommendations in Foreman-Mackey et al. (2012) and experimentation with the simulated datasets, we chose 16 walkers and selected random (valid) values for each walker’s starting point. After testing several combinations, we found that generally 300 burn-in steps and 100 chain steps per walker resulted in stable solutions. To aide in computational efficiency, we placed sensible restrictions on the allowed ranges for each of the three parameters (i.e., $p_{\text{prior}}(\theta)$ and $p_{\text{prior}}(N_{\text{pred}})$): $-6 \leq \alpha \leq 6$, $M_{\text{max,obs}} \leq M_{\text{max}} \leq 120 M_{\odot}$, and $0 < N_{\text{pred}} < 10^2 \times N$. In the case of finite mass uncertainties, we revised the restriction on M_{max} for reasons discussed in §4.1.

We refer the reader to the Appendix for a more in depth explanation of MCMC sampling theory.

3.2. Illustrative Examples with Idealized Data

We now illustrate this approach by applying it to idealized mock data (i.e., no mass uncertainties and a perfectly known boxcar completeness function). A simple way to create the simulated data for a power-law MF, assuming negligible mass uncertainties, is to define a variable x via $M = \exp(\ln[x(1-\alpha)]/(1-\alpha))$ and then draw a uniform random variate in x between $x_{\text{min/max,input}} = (1-\alpha)^{-1} M_{\text{max/min,input}}^{1-\alpha}$. For all illustrative examples in this paper we have adopted $\alpha_{\text{input}} = 2.35$ (i.e., Salpeter 1955). We discuss the results of simulations run with other values of α at the end of this section.

Each dataset consists of N masses that have been discretely sampled over the interval $[M_{\text{min,input}}, M_{\text{max,input}}]$. The value of $M_{\text{min,input}}$ is equivalent to the lower mass limit, M_{comp} , as specified by the completeness function. For these exercises, we have fixed $M_{\text{min,input}} = 3 M_{\odot}$, a conservative value for resolved clusters in M31 (e.g., Dalcanton et al. 2012).

The mass of the most massive star that could have formed in a cluster is of great interest astrophysically. Essentially it provides a clue to the fundamental nature of how stars form, e.g., is there a universal value for the most massive that can form or does it depend on local

environmental conditions? There are several studies in the literature that attempt to quantify the relationship between environment and the mass of the most massive star (e.g., Weidner et al. 2010; Cerviño et al. 2011; El-dridge 2012).

In our simulations, the value of the most massive star that could have formed in a cluster is $M_{\text{max,input}}$. The mass of the most massive star “observed”, i.e., on the mass list, is $M_{\text{max,obs}}$, which is less than $M_{\text{max,input}}$ due to the effects of stellar evolution or dynamical ejection. Under the assumption that the mass of each star is independently drawn from the IMF, it is fairly intuitive that $M_{\text{max,obs}}$ cannot inform us about $M_{\text{max,input}}$ beyond providing us a simple lower limit on $M_{\text{max,input}}$, i.e., the most massive observed stars does not contain information about more massive stars that are no longer in the cluster.

For a given dataset, we aim to recover the input MF slope, α_{input} , the maximal mass of a star that could have formed $M_{\text{max,input}}$, and require that our model produce a normalization that is consistent, i.e., within Poisson uncertainty, with the number of stars observed, N . We denote the corresponding model parameters as α , M_{max} , and N_{pred} .

We applied this technique to mock data for different permutations of α_{input} , $M_{\text{max,input}}$, and N . Results from the first example are shown in Figure 3. Here we show the recovered joint and marginalized distributions for a mock cluster which was generated assuming $N = 1000$ stars and $M_{\text{max,input}} = 60 M_{\odot}$. We recover $\alpha = 2.35^{+0.05}_{-0.05}$, $M_{\text{max}} = 54.0^{+2.6}_{-0.8}$, and $N_{\text{pred}} = 995^{+33}_{-31}$, where the indicated values represent the median and $0.5 \times (84\text{th} - 16\text{th percentile})$ of the corresponding marginalized distribu-




Figure 3. The recovery of MF parameters α , M_{max} , N_{pred} in the case of idealized mock data, i.e., negligible mass uncertainties and a boxcar completeness function, from a single run of `emcee`. The joint distributions are shown as a scatter plot in the center of the first three panels, for the indicated parameters. Dark colored points are close to the most probable value, while light colors are farther away. The solid black line encloses 68% of the points centered around the most probable combination of parameters. The corresponding marginalized distributions have been projected onto the axis of each plot. The solid red lines indicate the median of each marginalized distribution, and the shading encloses 68% of the distribution centered around the median. The roundness of each joint distribution indicates little degeneracy between the MF parameters. For this example dataset, we have recovered $\alpha = 2.35^{+0.05}_{-0.05}$, $M_{\text{max}} = 54.0^{+2.6}_{-0.8}$, and $N_{\text{pred}} = 995^{+33}_{-31}$. In the lower right hand panel, the observed cumulative mass function is shown as the thick purple line, and 100 MF models randomly drawn from the pPDF are over-plotted as thin grey lines.




Figure 4. The recovery of the MF parameters α , M_{\max} , N_{pred} in the case of idealized mock data, i.e., negligible mass uncertainties and a boxcar completeness function, from a single run of *emcee*. This simulation has $N = 500$ stars, $\alpha_{\text{input}} = 2.35$, and $M_{\max, \text{input}} = 60 M_{\odot}$. From top left to lower right: the marginalized distributions for α , M_{\max} , and N_{pred} . The solid red line represents the median of the distribution and the grey shaded region highlights 68% of the distribution centered around the median, i.e., the range enclosed by the 16th and 84th percentiles. In each panel, the blue dashed line represents the ‘truth’, i.e., α_{input} , $M_{\max, \text{obs}}$, and N , respectively. For this example dataset, we have recovered $\alpha = 2.35^{+0.07}_{-0.07}$, $M_{\max} = 63.3^{+9.0}_{-3.1}$, and $N_{\text{pred}} = 500^{+22}_{-21}$, indicating excellent consistency with the input values. In the case of M_{\max} , we are only able to place constraints on $M_{\max, \text{obs}}$ and not on the mass of the most massive star formed. In the lower right hand panel, the observed cumulative mass function is shown as the thick purple line, and 100 MF models randomly drawn from the pPDF are over-plotted as thin grey lines.

tion, i.e., $1-\sigma$ for a Gaussian distribution.

In general, the recovered values reflect the input values with excellent accuracy and precision. However, as expected, the recovered value of M_{\max} is consistent with the mass of most massive observed star, $M_{\max, \text{obs}}$, and not that of the most massive star that could have formed $M_{\max, \text{input}}$. Thus, this constraint provides a lower bound on the upper mass limit of the IMF.

The joint distributions in Figure 3 show little covariance between parameters, which indicates that there is little degeneracy in this problem, i.e., the estimation of one parameter will not affect the accuracy and precision to which another parameter can be constrained. Consequently, and for clarity, we will focus on the marginalized distributions for the next set of examples.

In Figures 4-6, we show select examples for tests of cluster MF recovery in the cases where $N = 500$, 100, and 25 stars and $M_{\max, \text{input}} = 60$, 30, and 15 M_{\odot} . The first three panels in each figure show the marginalized distributions for α , M_{\max} , and N_{pred} . For the case of 500 observed stars and $M_{\max, \text{obs}} = 59.9 M_{\odot}$ (Figure 4), we recover $\alpha = 2.35^{+0.07}_{-0.07}$, $M_{\max} = 63.3^{+9.0}_{-3.1}$, and $N_{\text{pred}} = 500^{+22}_{-21}$, which are all consistent with their input or observed (in the case of M_{\max}) values.

Figures 5 and 6 are examples of the same tests applied to populations with fewer stars and smaller dynamic ranges in mass. In both examples, the recovered parameters are consistent with the input values. As expected, as the amount of observational information decreases, the constraints become increasingly broad. However, even in the limit of 25 observed stars and $M_{\max, \text{input}} = 15 M_{\odot}$ (Figure 6), we find $\alpha = 2.50^{+0.60}_{-0.65}$ and $N_{\text{pred}} = 26^{+5}_{-5}$.




Figure 5. The same as Figure 4 only with $N = 100$ stars and $M_{\max, \text{input}} = 30 M_{\odot}$. Here, we recover $\alpha = 2.37^{+0.19}_{-0.20}$, $M_{\max} = 32.2^{+14.8}_{-3.5}$, $N_{\text{pred}} = 101^{+11}_{-10}$. The recovered fractional precision on α is roughly a factor of 2.5 larger than for the example considered in Figure 4.

Thus, these ranges contain the input/observed values for each parameter.

In general, these exercises demonstrate probabilistic approach in §2 can recover MF parameters, even with limited information. In other words, useful constraints can be derived even for clusters that fail to reach a ‘gold-standard’ level. Importantly, we verify that the recovered MF slope is equivalent to the input IMF slope, as is expected in the case of a single age cluster (e.g., Elmegreen & Scalo 2006). This finding affirms that this probabilistic technique provides the means to directly infer the slope of the high mass IMF of a given cluster with full and accurate accounting of the associated uncertainties, in the absence of dynamical effects, unresolved binaries, etc., as discussed in §6.

On the other hand, the recovered constraints on M_{\max} are less informative about the most massive star that could have formed, i.e., the upper mass limit of the IMF. We find that the recovered values of M_{\max} only constrain $M_{\max, \text{obs}}$, which is only a lower limit on $M_{\max, \text{input}}$, the IMF maximum mass limit for a given cluster. We there-




Figure 6. The same as Figure 4 only with $N = 25$ stars and $M_{\max, \text{input}} = 15 M_{\odot}$. Here, we recover $\alpha = 2.50^{+0.60}_{-0.65}$, $M_{\max} = 13.6^{+25.2}_{-2.3}$, $N_{\text{pred}} = 26^{+5}_{-5}$. The recovered fractional precision on α is roughly a factor of 8 larger than for the example considered in Figure 4. This exercise demonstrates that it is still possible to place infer the slope of the MF from a sparsely populated cluster.




Figure 7. The accuracy and precision to which we can recover α , in the case of negligible mass uncertainties and a boxcar completeness function. *Left panel* – The distribution of median values of α from 100 recovered datasets, each with 10^4 stars and $M_{\text{max,input}} = 60 M_{\odot}$. The resulting distribution is well-characterized by a normal distribution with $\mu = 0.006$ and $\sigma = 0.07$, indicating that this approach to MF recovery is free of systematic biases. *Right panel* – A 2D map of the theoretical precision to which α can be recovered as a function of the number of stars and dynamic mass range. This plot indicates that that intrinsic precision (i.e., $0.5 \times (84\text{th}-16\text{th percentiles})$) for the recovery of α ranges from < 0.1 in the limit of many stars and a high dynamic mass range to > 0.7 for a poorly populated cluster with a low dynamic mass range. However, under the assumption that masses are independently drawn from the MF, it is possible to combine multiple small clusters to improve precision of MF recover. The asymmetry in the distribution indicates a slight increase in the recovery of α for a small cluster with a high dynamic mass range relative to a well-populated cluster with a low dynamic mass range.

fore suggest that dedicated searches for the most massive star (e.g., spectroscopic surveys) are more suited to constraining the upper stellar mass limit of the IMF than application of this technique (e.g., Massey et al. 1995a; Lennon 1997; Massey 2003; Kudritzki et al. 2008; Crowther et al. 2010; Evans et al. 2011). Consequently, for the remainder of the paper, we will primarily focus on the recovery of α , and only discuss constraints on M_{max} in cases of particular interest.

Extending the above exercises to larger ranges of N and M_{max} allows us to illustrate the precision and accuracy to which α can be recovered for a diverse set of clusters. We first check the accuracy of the method to verify that there are no systematic biases in the recovered MF slope. To do this, we simulated 100 different clusters each with 10^4 stars, to minimize stochastic effects, and $M_{\text{max,input}} = 60 M_{\odot}$, for the same underlying MF. We recovered the MF for each of the 100 clusters and plot the distribution of the difference between α_{input} and the median value of $\alpha_{\text{recovered}}$ in the left panel of Figure 7. The mean of the distribution is 0.006 and $0.5 \times 84\text{th} - 16\text{th percentiles}$, i.e., a Gaussian $1-\sigma$, is ± 0.07 . Based on this set of tests, we see no evidence for systematic biases in the recovery of α in any of the parameter space considered.

Next, we explore the theoretical precision to which α is recovered as a function of N and the dynamic range of the observed masses, as proxied by $\log(M_{\text{max,obs}}/M_{\text{comp}})$. We constructed a grid of model clusters spanning a range in N and $\log(M_{\text{max,obs}}/M_{\text{comp}})$ as follows. We fixed $M_{\text{comp}} = 3 M_{\odot}$ and created 7000 simulated clusters by drawing random values N and $M_{\text{max,input}}$ such that $0.8 \leq \log(N) \leq 3.5$ and $0.2 \leq \log(M_{\text{max,obs}}) \leq 1.3$, and discretely sampled the IMFs following the procedure outlined at the beginning of this section. We recovered the MF slope for each realization and computed $\Delta\alpha = 0.5 \times (84\text{th} - 16\text{th percentile})$, i.e., a percentile based $1-\sigma$ equivalent. To account for fluctuations due to the discrete na-

ture of the sampling, we applied a Gaussian smoothing kernel to the resulting grid with a $1-\sigma$ width of $\sim 5\%$ of the dynamic range in each variable. We computed the smoothed and un-smoothed values and found they differ by no more than $\sim 7\%$ at any point.

We have plotted the results of this exercise in the right panel of Figure 7. In the limit of high observational information, e.g., $N \sim 10^3$ stars and $\log(M_{\text{max,obs}}/M_{\text{comp}}) \sim 1.1$, we see that α is recovered with a precision of $\sim \pm 0.1$. At the other extreme ($N \sim 10$ stars and $\log(M_{\text{max,obs}}/M_{\text{comp}}) \sim 0.3$), the precision in $\alpha_{\text{recovered}}$ decreases significantly to $\gtrsim \pm 0.7$. Because the mock data used for these tests are highly idealized, i.e., perfectly known individual stellar masses and completeness, these values reflect the highest theoretical precision with which α can be inferred from a given number of stars and a measured dynamic mass range. As we discuss in the next sections, other considerations such as finite mass uncertainties and completeness will degrade these levels of precision, lessening the achievable constraints.

Additionally, using a least squares fitting routine, we find that $\Delta\alpha$ can be reasonably approximated as a function of $\log(N)$ and $\log(M_{\text{max,obs}}/M_{\text{comp}})$ by a 3rd order polynomial

$$\Delta\alpha = \sum_{i=0}^3 \sum_{j=0}^3 \beta_{i,j} (\log(N))^i (\log(M_{\text{max,obs}}/M_{\text{comp}}))^j, \quad (16)$$

where the coefficients $\beta_{i,j}$ are listed in Table 2. Within the range of $\log(N) = 1.5$ to 3.0 and $\log(M_{\text{max,obs}}/M_{\text{comp}}) = 0.4$ to 1.2, this approximation is generally good to within $\sim 20\%$. Outside this range, this analytic expression is noticeably steeper, meaning the extrapolations into regimes of lower and higher information content lead to over- and underestimates of the true precision. Nevertheless, when used with a reasonable level of caution, this analytic expression provides




Figure 8. *Left* – The 2D map of the theoretical precision to which α can be recovered from Figure 7. We have over plotted the indicated precision from the literature studies shown in Figure 1 and tabulated in Table 1. Points that appear darker than the surrounding regions have listed errors above the theoretical precision limits, where as lighter points have errors that are smaller than are theoretically permissible. *Right* – A histogram showing the distribution of the ratio of quoted literature errors and our computed theoretical limits for clusters show in Figure 8 and tabulated in Table 1. The red shaded region indicates the region in which literature studies have quoted errors less than the theoretical limit. The range of ratios varies over an order of magnitude, and approximately $\sim 3/4$ of the literature studies quote error bars than are below theoretical expectation.

a good rule of thumb for the best attainable precision on the MF slope.

Finally, we tested the accuracy and precision to which other α values can be recovered. Differing values of α dictate the relative distribution of stellar masses, and a significantly steeper MF slope would produce fewer massive stars, potentially resulting in less accurate and/or precise constraints on the MF slope. To test this potential dependence on α , we selected two different, but plausible values of $\alpha = 1.8$ and 2.8 , and repeated all the simulations presented in Figure 7. We found that both the accuracy and precision to which α can be recovered is statistically consistent in all cases. There were random variations at the few percent level; however, such small differences are easily accounted for by the discrete nature of the sampling and the finite number of realizations.

3.3. Comparing Theoretical Precision Limits with Values from the Literature

As an informative exercise, in the left panel of Figure 8 we have over-plotted the error bars of the literature MF values from Figure 1 on top of our derived theoretically attainable values. The contrast in colors between the points and shaded contours indicates the level of agreement between the two; points that are lighter than the surrounding region are below the theoretically attainable precision, while darker points are larger than this lower limit. A cursory inspection of this figure indicates a full range of differences between the two datasets, i.e., there are a number of points with extreme light and dark color contrast. We have tabulated the literature data used to construct this plot in Table 1.

To better quantify differences in the two datasets, in the right panel of Figure 8 we have plotted a histogram of the ratios between the theoretical and literature MF precision values. The histogram shows that the ratios vary over an order of magnitude around unity, with $\sim 3/4$ of the error bars quoted in the selected literature studies being smaller than the theoretically attainable limit (red-shaded region). The underestimated error bars are

typically lower than the theoretical expectations by a factor of ~ 2 or more. Overly small error estimates can give the false impression of a good constraint, which could lead to claims of a significantly unusual MF measurement. Conversely, large overestimates of the errors on a MF slope indicate that maximal information may not be extracted from the data. Both cases reinforce the importance of developing a well-defined methodology for correctly characterizing the MF slope of a young cluster.

Using the literature measurements of α in tandem with the theoretical lower limit errors, we compute a mean α for the literature values considered. In Figure 9, we




Figure 9. A modified version of Figure 1 based on a comparison between the amplitude of the error bars quoted in the literature and the theoretical lower limits derived in this paper. The blue points are studies whose quoted errors are larger than the theoretical lower limits. The magenta points are those that had quoted error bars smaller than the theoretical lower limit. For these points, we have assigned and plotted new error bars equal to the theoretical lower limits. We then use all the literature data to compute the weighted mean (solid green line) and weighted $1-\sigma$ (lightly shaded green region) of α , which we find to be $\langle \alpha \rangle = 2.46$ with a $1-\sigma$ dispersion of 0.35 dex. Note that while we are able to assign more robust error bars to most points, we are unable to assess the accuracy of each value of α , which would require re-analysis of each dataset.

have re-created Figure 1, but have plotted the theoretical lower limits errors for those studies which have underestimated the uncertainties. These are indicated as magenta points. Studies whose error bars are larger than the theoretical limit have been left as is, and are plotted as blue points. From the Figure 9 and Table 1, we see there is no discernible trend for which studies, values of α , or mass ranges have underestimated error bars. Using all reported α values, along with the new set of uncertainties, we computed a weighted mean and weighted 1- σ uncertainty for this selection of IMF studies and find $\langle\alpha\rangle = 2.46$ with a 1- σ dispersion of 0.35 dex. This value is consistent with the slope of a Salpeter and Kroupa IMFs, $\alpha = 2.35$ and 2.3, respectively. However, the scatter is large ensuring the mean value is also consistent with the Kennicutt IMF (Kennicutt 1983), $\alpha = 2.5$, and the Scalo IMF (Scalo 1986), $\alpha \sim 2.6$. Although our computed mean value of α depends on our selection of literature studies and the accuracy of each study, the broad scatter reinforces the notion that our knowledge of the high mass IMF slope is far from secure.

4. INCORPORATING UNCERTAINTIES IN STELLAR MASSES

Up to this point, we have explored placing constraints on MF parameters presuming data with negligible individual mass uncertainties. However, in practice, the mass of a star is not perfectly known. Degeneracies with extinction, metallicity, etc., along with uncertainties in stellar models all contribute to uncertainties in estimating the mass of a star.

Clearly, the uncertainty on the mass of an a star needs to be incorporated into the pPDF for realistic MF determination. To understand the relationship between MF determination and stellar mass uncertainties, it is useful to review how stellar masses are measured.

Most individual star masses are inferred from photometric spectral energy distribution (SED) fitting, i.e., by




Figure 10. A schematic demonstration of the importance of considering the MF as a prior in the determination of individual stellar masses. The result of photometric SED fitting for a stellar masses, i.e., $p(\vec{d} | M)$, is denoted by the blue line. The MF with $\alpha = 2.35$ (red dot-dashed line), indicates that lower stellar masses are more likely than higher stellar masses. Therefore, the best mass estimate for the star comes from the convolution of these two functions (green dashed line). Not taking the MF into account when deriving stellar mass PDFs can lead to noticeable overestimates of the star’s true mass.

comparing photometric observations of a star, \vec{d} , with predicted fluxes from stellar models. Specifically, one can construct a set of predicted fluxes in the observed bands by for various combinations of $\log(T_{eff})$, $\log(g)$, Z , A_V , M , age, etc., and then perform a comparison, e.g., χ^2 -minimization, between the observed and predicted fluxes. From this process one is left with a multi-dimensional distribution for the probability of the star’s parameters given the data, i.e., $p(\vec{d}_i | \log(T_{eff}), \log(g), Z, A_V, M, \text{age})$. One can then compute the likelihood of the data given a mass, i.e., $p(\vec{d} | M)$, by marginalizing over all other model parameters. The resulting PDF for a star’s mass accounts for correlations and uncertainties in the all other parameters.

While the quantity $p(\vec{d} | M)$ is often thought to reflect the probability of a star’s mass, the desired quantity is in fact, $p(M | \vec{d})$, the probability that a star has a particular mass given the observations. To determine this quantity, one can apply Bayes’s theorem, which yields

$$p(M | \vec{d}) = \frac{1}{\zeta} p(\vec{d} | M) p_{\text{prior}}(M), \quad (17)$$

where $p_{\text{prior}}(M)$ reflects any prior knowledge about the probability of individual stellar masses, and ζ is a normalization factor needed to make $p(M | \vec{d})$ a true probability.

A temptingly simple assumption for $p_{\text{prior}}(M)$ would be that all stellar masses are equally likely, i.e., $p_{\text{prior}}(M) = \text{constant}$. In this case, $p(M | \vec{d})$ and $p(\vec{d} | M)$ can be treated interchangeably, which is at least implicitly assumed in many published analyses (e.g., Massey et al. 1995a,b; Panagia et al. 2000; Brandner et al. 2008; Bianchi et al. 2012).

However, it is clearly more appropriate to adopt the cluster’s MF as a prior on the distribution of individual stellar masses. As illustrated in Figure 10, failure to consider the effects of the MF when determining the mass of a star will lead to a biased mass relative to the star’s true mass in the presence of significant mass errors.

Adopting the MF from Equation 4 as a prior on stellar mass, we can now write down an expression for the probability of the data, \vec{d} , given MF parameter, $\vec{\theta}$, as:

$$\begin{aligned} p(\vec{d} | \vec{\theta}, \text{obs}) &= \int p(\vec{d} | M) p_{\text{MF}_0}(M | \vec{\theta}, \text{obs}) dM \\ &= \frac{p(\text{obs} | \vec{d}) \int p(\vec{d}_i | M) p_{\text{MF}}(M | \vec{\theta}) dM}{p(\text{obs} | \vec{\theta})}, \end{aligned} \quad (18)$$

and

$$p(\text{obs} | \vec{\theta}) = \int p(\text{obs} | \vec{d}) p(\vec{d} | \vec{\theta}) d\vec{d}, \quad (19)$$

where $p(\vec{d} | M)$ is the likelihood of the observed data given a true mass, as determined by photometric SED fitting, and $p(\text{obs} | \vec{\theta})$ is the normalization factor, which represents the probability of observing the data (i.e., fluxes) given the MF model parameters $\vec{\theta}$.

Using Bayes’s theorem, we can write the probability for the parameters, $\vec{\theta}$ and N_{pred} , given the observed data, \vec{d} as:

$$p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \vec{d}, \text{obs}) = p_{\text{poi}}(N | N_{\text{pred}}) \times \sum_{i=1}^N \frac{p(\text{obs} | \vec{d}_i) \int p(\vec{d}_i | M) p_{\text{MF}}(M | \vec{\theta}) dM}{p(\text{obs} | \vec{\theta})} p_{\text{prior}}(\vec{\theta}) p_{\text{prior}}(N_{\text{pred}}), \quad (20)$$

which is the pPDF in the case of non-negligible mass uncertainties.

4.1. Illustrating the Effects of Log-Normal Mass Uncertainties

As an example of how to fold in finite mass uncertainties in practice, we consider the case that each star has an observed mass PDF that can be represented by a log-normal function. The assumption of a log-normal not only provides an analytically tractable solution, which aides in the clarity of this paper, but also log-normal approximations (i.e., $\log(M) \pm \delta m$) have been used in the literature to summarize the inferred masses of stars (e.g., Massey & Hunter 1998; Massey 2011; Bianchi et al. 2012).

Consider a set of N stars each of which has an associated mass PDF. We believe that each observed mass PDF is the convolution of a single true mass, M_i , with a log-normal noise model (due to observational uncertainties, the flux-to-mass conversion process, etc.). The inferred mass for each star can then be summarized as having a mean value of \bar{M}_i and a fractional mass uncertainty, $f_i \equiv M_i/\bar{M}_i$, i.e., the linear width of the mass PDF scales as $\sim e^f$. Considering only the upper portion of the IMF, the stars we observe can have true masses ranging from M_{min} to M_{max} , while the inferred masses will range according to the stellar evolution models and the observed completeness function, and thus have limits of M_{comp} and M_{lim} . For simplicity, we will adopt a boxcar completeness function written in terms of \bar{M} ,

$$p(\text{obs} | \bar{M}) = \begin{cases} 1, & M_{\text{comp}} \leq \bar{M} \leq M_{\text{lim}} M_{\odot} \\ 0, & \text{otherwise} \end{cases} \quad (21)$$

Having defined the relevant variables and the ranges, we can return to the derivation. The goal of this process is to place constraints on the MF parameters, given the set of log-normal PDFs for the inferred masses, and Equation 18 provides the framework to complete this task. Substituting a log-normal mass model for $p(\vec{d}_i | M)$ and making the coordinate transformation $X \equiv \ln M/\bar{M}$, allows us to write:

$$p(\vec{d}_i | \vec{\theta}, \text{obs}) = c_{\text{MF}_0}(\vec{\theta}) \bar{M}^{-\alpha} \frac{p(\text{obs} | \bar{M})}{\sqrt{2\pi}f} \int_{\ln(\frac{M_{\text{min}}}{\bar{M}})}^{\ln(\frac{M_{\text{max}}}{\bar{M}})} e^{-\frac{X^2}{2f^2}} e^{(1-\alpha)X} dX, \quad (22)$$

where the completeness is over the observed masses, $p(\text{obs} | \bar{M})$, the integration is over the permissible range of true masses, and we have made the simplifying assumption of a constant fractional mass uncertainty for all masses, i.e., $f_i \rightarrow f$. This assumption has been made for clarity in this derivation, but in practice, f_i is not

likely to be constant. Stars with $M \gtrsim 25 M_{\odot}$ are less constrained than those of lower mass stars (e.g., Massey et al. 1995a), owing to both uncertainties in massive star evolution models, as well as degeneracies in the optical and near-UV photometric colors of massive stars, which make it difficult to precisely characterize extremely massive stars.

Integration of Equation 22 yields a closed form solution:

$$p(\vec{d} | \vec{\theta}, \text{obs}) = c_D(\vec{\theta}) c_{\text{MF}_0}(\vec{\theta}) \bar{M}^{-\alpha} p(\text{obs} | \bar{M}) e^{\frac{(\alpha-1)^2 f^2}{2}} \times \left(\eta \left[\frac{(\alpha-1)f^2 + \ln(\frac{M_{\text{max}}}{\bar{M}})}{f} \right] - \eta \left[\frac{(\alpha-1)f^2 + \ln(\frac{M_{\text{min}}}{\bar{M}})}{f} \right] \right), \quad (23)$$

where

$$c_D(\vec{\theta})^{-1} = c_{\text{MF}_0}(\vec{\theta}) e^{\frac{(\alpha-1)^2 f^2}{2}} \int_{M_{\text{comp}}}^{M_{\text{lim}}} \bar{M}^{-\alpha} p(\text{obs} | \bar{M}) \times \left(\eta \left[\frac{(\alpha-1)f^2 + \ln(\frac{M_{\text{max}}}{\bar{M}})}{f} \right] - \eta \left[\frac{(\alpha-1)f^2 + \ln(\frac{M_{\text{min}}}{\bar{M}})}{f} \right] \right) d\bar{M}, \quad (24)$$

and η is the normalized cumulative distribution function of a standard Gaussian PDF with $\mu = 0$ and $\sigma = 1$:

$$\eta(X) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^X e^{-\frac{x^2}{2}} dx. \quad (25)$$

In the limit of very accurate mass estimates, i.e., $f \rightarrow 0$, the expression for $p(\vec{d}_i | \vec{\theta})$ reduces to a scenario where the masses are perfectly known, as in Equation 13.

Having derived an expression for $p(\vec{d}_i | \vec{\theta})$, we can now write the pPDF for the MF parameter as:

$$p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \{\vec{d}_i\}, N, \text{obs}) = p_{\text{poi}}(N | N_{\text{pred}}) \sum_{i=1}^N \ln p(\vec{d}_i | \vec{\theta}, \text{obs}) p_{\text{prior}}(\vec{\theta}) p_{\text{prior}}(N_{\text{pred}}), \quad (26)$$

4.2. Illustrative Examples with Mass Uncertainties

In this section, we present several examples of MF recovery in the presence of non-negligible mass uncertainties. However, before delving into specific cases, it is instructive to first examine how log-normal mass PDFs affect the observed MF.

4.2.1. Schematic Examples with Continuous Mass Functions

In Figure 11, we illustrate the changes in the (continuous) observed MF for different values of f . For reference, the solid black line represents a Salpeter IMF from 3 to 120 M_{\odot} . The other plotted lines represent MFs generated using Equation 22 for select values of f (0.1, 0.5,




Figure 11. A schematic demonstrating the influence of log-normal individual stellar mass PDFs on the MF. Each distribution was constructed assuming $M_{\max, \text{input}} = 40 M_{\odot}$ (grey line) and $\alpha_{\text{input}} = 2.35$ (black line). We then varied the value the fractional mass uncertainties, f , as indicated. Small fractional mass uncertainties lead to a slightly steeper than Salpeter MF slope, and predicted the presence of masses larger than $M_{\max, \text{input}}$. For extreme values of f , e.g., 1, the observed MF slope is flatter than Salpeter, with observed values of M_{\max} significantly larger than $M_{\max, \text{input}}$. Failure to properly account for mass uncertainties when modeling the observed MF can lead to biases in the recovery of α and M_{\max} . See §4.1 for more discussion.

1.0, 2.0). In each case, $M_{\max, \text{input}} = 40 M_{\odot}$ and each MF was normalized to have the same number of $3 M_{\odot}$ stars as a Salpeter MF.

For the smallest fractional mass uncertainty, $f = 0.1$, we see that the observed MF (solid blue line) is identical to Salpeter until $\sim 37 M_{\odot}$. At this point it begins to deviate from the Salpeter MF, and predicts the presence of stars above the nominal value of $M_{\max, \text{input}} = 40 M_{\odot}$. Taken at face value, this MF appears slightly steeper and it implies a finite probability that $M_{\max, \text{obs}}$ is larger than $M_{\max, \text{input}}$. The MF for $f = 0.5$ exhibits the same qualitative behavior, with an even steeper slope and a significantly larger apparent value of $M_{\max, \text{obs}}$. Naive fitting of the observed MF, without taking mass errors into account, would therefore lead to biased measurements of both α and M_{\max} .

For the extreme values of $f=1$ and 2, the observed MFs flatten out relative to Salpeter, and each predicts a significant population of stars with masses in excess of $100 M_{\odot}$. For such extreme uncertainties, a significant number of stars with true masses below $3 M_{\odot}$ will have inferred masses above the completeness limit.

This schematic also illustrates the need to alter the prior restrictions placed on M_{\max} . In the case of no mass uncertainties, the inferred and true masses are identical, implying that $M_{\max} \geq M_{\max, \text{obs}}$ was a reasonable requirement. However, for finite mass uncertainties, $M_{\max, \text{obs}}$ is likely an overestimate of $M_{\max, \text{input}}$, i.e., the maximal inferred mass is larger than the maximal true mass, invalidating the previous assumption. For the exercises involving finite mass uncertainties, we therefore only require that $M_{\max} \leq M_{\text{lim}}$, the larger mass permitted by stellar evolution models.

4.2.2. Illustrative Examples with Simulated Clusters

We now use mock data to explore recovery of the MF in the case that each mass has a log-normal PDF. To generate the simulated data, we construct a convolution

between the MF a log-normal noise model with $\alpha = 2.35$ (Equation 22). We then draw 10^6 stars from this function, apply the boxcar completeness function from Equation 21, and randomly subsample the mass list to obtain the desired number of stars.

In Figures 12 and 13, we show results from the recovery of α for select simulated clusters. In each case, the mock data consisted of 10^4 stars (to minimize random noise effects) and discrete values of $f = 0.1, 0.5$, and 1.0. To illustrate the effects of the observed dynamic mass range, we also considered two different values of $M_{\max, \text{input}} = 20 M_{\odot}$ and $60 M_{\odot}$. As a baseline for comparison, we applied two versions of the code to each dataset, one that included modeling of the mass uncertainties (the ‘noisy’ MF model; Equation 26) and the simple model MF (the ‘noiseless’ MF model; Equation 15). In each panel, the dashed line indicates the input value of α while the solid line indicates the median value for each marginalized PDF.

In Figure 12, we show the recovery of α for $M_{\max, \text{input}} = 60 M_{\odot}$. In the case of small mass errors ($f = 0.1$), both the noisy (red) and noiseless (grey) model MF models provide excellent recovery of α_{input} . The small level of offset between the results are consistent with the expected scatter ($\sim \pm 0.03$) in the case of $N = 10^4$ stars. The two methods also return the same level of precision on $\alpha_{\text{recovered}}$, as indicated by the comparable widths of the two distributions.




Figure 12. The recovered marginalized distributions for α when accounting (colored histograms) and not accounting (grey histograms) for individual stellar mass uncertainties in the MF model (colored histograms) and not including mass uncertainties (grey histograms). Each star has a log-normal mass PDF with a fractional uncertainty, $f \equiv M/M$, whose assumed constant value is indicated. For each value of f the noisy (i.e., accounting for mass uncertainties) and noiseless (i.e., not accounting for mass uncertainties) models were applied to the same data, which always had $M_{\max, \text{input}} = 60 M_{\odot}$ and $N = 10^4$ stars. In each of the first three panels, the true MF slope (Salpeter), is indicated by the dashed black line. Application of the noisy model yields recovery of the colored histograms, whose median value is indicated by the solid colored line. The grey histograms are the result of applying the noiseless models, and the solid black line represents the median of this distribution. For reasonable fractional mass uncertainties, e.g., $f \lesssim 0.5$, there are only small differences in the MF slopes recovered by the two methods. For the extreme case, $f = 1.0$, modeling the mass uncertainties yields near perfect recovery, while failure to do so results in a systematically flatter MF slope. In the lower right panel, we show the MF distributions used in each case, and have included a Salpeter MF (with an arbitrary vertical scaling applied) for reference.




Figure 13. The same as Figure 12 only with $M_{\text{max,input}} = 20 M_{\odot}$. In this case, the differences in the recovered MF slopes between the two methods are more pronounced, even for modest values of f . These examples demonstrate the importance of including a good characterization of mass uncertainties in order to measure the MF slope.

For significantly larger mass errors ($f = 0.5$), both the noisy (green) and noiseless (grey) models still provide excellent recovery of α . The distributions again share a comparable width. Although the noiseless PDF appears shifted toward slightly higher values of α , the median recovered value is within the expected scatter for a cluster with $N = 10^4$ stars.

Extending this exercise to extremely large mass errors ($f = 1.0$) reveals significant discrepancies between results from the noisy (blue) and noiseless (grey) models. Based on the schematic in Figure 11, the observed MF should be flatter than Salpeter, which is what is the noiseless model returns, with the median $\alpha_{\text{recovered}} = 2.20$. Although the distribution is comparable in width to those in the other panels of the plot, it does not overlap the true value of α . In contrast, application of the noisy model recovers α to excellent accuracy, although to a lower degree of precision.

To test the sensitivity of MF slope recovery to M_{max} , we conducted the same exercise, only with $M_{\text{max,input}} = 20 M_{\odot}$. The recovered distributions for α are shown in Figure 13. For the lower value of $M_{\text{max,input}}$, we see more pronounced differences in the results from the noisy and noiseless models. For both $f = 0.1$ and 0.5 , the noiseless model recovers systematically higher values of α relative to the input. While only a modest bias is present in the case of $f = 0.1$, the offset is quite substantial when $f = 0.5$. In contrast, the recovered values of α from application of the noisy MF model are in excellent agreement with α_{input} . Once again, the widths of the noisy model distributions are slightly broader than the noiseless model results. However, the slight decrease in precision is a small tradeoff relative to inferring a significantly more accurate MF slope.

In the case of $f = 1.0$, application of the noiseless MF model to the data yields an only slightly biased ($\lesssim 0.1$) median value of α . This is perhaps somewhat surprising, given the extremely large fractional uncertainties. However, this small amount of disparity is purely coincidental. The combination of f and $M_{\text{max,input}}$ conspired to produce an observed MF slope that appears close to the true MF slope. Such a case might lead to the conclu-

sion of a precise constraint on a MF that is statistically different from Salpeter, despite the large mass errors; however, such a conclusion would clearly be false. In contrast, modeling the mass uncertainties results in an accurate recovery of the true MF, albeit with decreased precision, as previously discussed.

Overall, we find that incorporating the effects of mass uncertainties into MF modeling results in an unbiased recovery of the MF. In contrast, failure to model even moderately large mass uncertainties, e.g., $f \gtrsim 0.5$, can lead to biases that are difficult to quantify, which can undermine a physically meaningful interpretation of the results. Additionally, neglecting to model mass uncertainties can lead to a severely low estimate of M_{max} , which is equivalent of underestimating the cluster's age. Finally, as previously mentioned, in practice f_i is unlikely to be constant, as the inferred properties of stars with $M \gtrsim 25 M_{\odot}$ are far more uncertain than lower mass stars (e.g., Massey et al. 1995a). Including a non-constant value of f is straightforward in the presented framework.

5. INCORPORATING THE EFFECTS OF OBSERVATIONAL COMPLETENESS

5.1. Preface

We now consider the effects of observational completeness on the inference of MF parameters. For clarity in the mathematics and the examples, we return to the case of perfectly known masses. In practice, both completeness and mass uncertainties need to be simultaneously modeled, but they typically affect opposite ends of the observed mass spectrum, meaning that they are largely not degenerate in terms of the data they impact.

When deriving the general pPDF for a cluster's MF (Equation 11), we implicitly included the completeness function in the probability of a star being on our data list, $p_{\text{MF}_o}(M | \vec{\theta}, \text{obs})$. However, writing the pPDF to explicitly include the effects of completeness provides a better illustration of the role it plays. Using Equation 4, we can write the pPDF for a general completeness function as:

$$p_{\text{post}}(\vec{\theta}, N_{\text{pred}} | \{M_i\}, N, \text{obs}) = p_{\text{poi}}(N | N_{\text{pred}}) \times \sum_{i=1}^N (p(\text{obs} | M) c_{\text{MF}_o}(\vec{\theta}) M_i^{-\alpha}) p_{\text{prior}}(\vec{\theta}) p_{\text{prior}}(N_{\text{pred}}), \quad (27)$$

Here we see that the completeness enters into the pPDF as a linear weight on the MF. In the limit that the $p(\text{obs} | M) = \text{constant}$, it will have no effect on the pPDF as all points are equally weighted. In the next section we illustrate how a more realistic completeness function affects the recovery of the MF.

5.2. Illustrating the Effects of a Linear Completeness Function

In practice, a boxcar function is not a realistic representation of observational completeness. Completeness functions in clusters vary widely in form depending largely on the surface brightness of the cluster and the resolution of the observations. In this section, we consider a linear ramp completeness function. Such a linear completeness function is analytically tractable and provides a reasonable first order approximation to more realistic completeness functions presented in the literature (e.g., Anderson et al. 2008; Gennaro et al. 2011).

The linear ramp function, which ranges from 0 to 100% completeness, can be written piecewise as

$$p(\text{obs} | M) = \begin{cases} 1, & M > M_{\text{Cmax}} \\ 0, & M < M_{\text{Cmin}} \\ \frac{M - M_{\text{Cmin}}}{M_{\text{Cmax}} - M_{\text{Cmin}}}, & M_{\text{Cmin}} \leq M \leq M_{\text{Cmax}} \end{cases}, \quad (28)$$

where M_{Cmax} and M_{Cmin} are the upper and lower bounds of the linear portion of the completeness function.

We again turn to mock data with no mass uncertainties to illustrate MF recovery in the case of a linear completeness function. Specifically, we consider the cases of how perfectly known, slightly incorrect, and extremely incorrect knowledge of the completeness function affects MF recovery.

To generate the mock data we followed the general description in §3.2, with a few modifications. After drawing a list of masses from the power-law MF, we applied the ‘true’ completeness function to the mass list. Motivated by data from PHAT, for this exercise, the true completeness function is always a linear ramp between $M_{\text{Cmin}} = 2$ and $M_{\text{Cmax}} = 4 M_{\odot}$, with a completeness fraction of 0 for stars with $M \leq 2 M_{\odot}$ and 1 for $M \geq 4 M_{\odot}$. After applying this completeness function to the mass list, we then rejected stars from the list with a probability equal to their completeness fraction, e.g., a mass with $p(\text{obs} | M) = 0.3$ has a 30% chance of being observed. We then perform a uniform draw from the remaining stars to arrive at the desired number of stars. This set of masses constitutes the true mass list. To examine effects of a non-perfectly known completeness function, we followed the same procedure as described above, only the values of M_{Cmin} and M_{Cmax} are different when applied to the mock data and when used to model the completeness in the MF recovery.

Keeping the same true completeness function, we then explored MF recovery for various assumed completeness functions. In Figure 14, we considered the case in which the true (grey) and assumed (red) completeness functions were identical linear ramps between 2 and 4 M_{\odot} . We simulated 100 datasets each with 10^5 stars and $M_{\text{max,input}} = 60 M_{\odot}$, and measured the median value of $\alpha_{\text{recovered}}$, whose distribution is plotted as the red histogram in the right panel. The median of this distribution is represented by the solid red line. For a perfectly known completeness function, we found excellent recovery of the MF slope, as expected.

We next considered the cases where the true and assumed completeness functions were not identical. In the first case, the ramp in the assumed completeness function occurred between $M_{\text{Cmin}} = 2$ and $M_{\text{Cmax}} = 3 M_{\odot}$ (green), meaning the data are assumed to be more complete actually are, i.e., the stars between 2 and 4 M_{\odot} are assigned artificially high weights in Equation 27. As before, we simulated 100 different clusters of 10^5 stars, and examined the distribution of the median recovered values of α . As shown in the right-hand panel of Figure 14, the recovered values of MF slope are clustered around $\alpha = 2.0$, with small scatter. This test indicates that a moderate overestimation of the true completeness function can induce a significant underestimate of α , because stars that are ‘missing’ from the lower mass end due to ob-




Figure 14. The recovery of α from various linear completeness functions. *Left Panel* – Modeled linear completeness functions that are perfect (red dotted), moderately underestimated (blue dashed), and moderately overestimated, compared to the ‘true’ completeness of the data (grey). In this example the true completeness is a linear from 2 to 4 M_{\odot} , with a value of 0% below 2 M_{\odot} and 100% above 4 M_{\odot} . *Center Panel* – A histogram of the data (grey) with the ‘true’ completeness function applied. The case of perfectly modeled completeness is in red. A moderate underestimate of the completeness (blue) indicates that more lower stars are present than expected, whereas a moderate overestimate of the true completeness (green) indicates a deficit of lower mass stars relative to observations. *Right Panel* – A series of histograms indicating the recovered median values of α for each of the modeled completeness functions. The distribution of median values is derived from 100 different datasets, each with 10^5 stars and $M_{\text{max,input}} = 60 M_{\odot}$. The solid colored lines indicate the median of each distribution, and grey-dashed line is the α_{input} . A well-characterized completeness function (red) results in an accurate recovery of α , whereas a modest overestimate (green) or underestimate (blue) lead to systematic under and overestimates of α .

servational completeness are instead interpreted as being underproduced by a flatter MF. This finding is similar to that in Ascenso et al. (2009), who demonstrates that an artificial flattening of the MF due to completeness considerations mimics the behavior of predicted mass segregation effects. We discuss this point further in §6.

We also considered a case of a modest underestimation of the true completeness function. Here, the mass limits on the assumed completeness function were $M_{\text{Cmin}} = 2 M_{\odot}$ and $M_{\text{Cmax}} = 6 M_{\odot}$. In this scenario, stars between 2 and 6 M_{\odot} are assigned artificially low weights in Equation 27. Application of this assumed completeness function results in median recovered values of $\alpha = 2.7$ (blue in Figure 14). This exercise indicates that even a moderate underestimation of the true completeness function can induce a significant overestimate of α .

Some MF studies in the literature either apply a conservative completeness limit, i.e., only consider data that are likely to be 100% complete, or make no attempt to correct for completeness. We explore the effects of each of these scenarios and present the results in Figure 15. In the case that no attempt is made to account for completeness (green), i.e., all data above 3 M_{\odot} are complete, the recovered slope of the MF is ~ 1.7 , which is significantly flatter than the true MF slope of 2.35.

The scenario in which a conservative completeness cut is applied is more promising (blue in Figure 15). Here, we see that the input value of α is accurately recovered, albeit with lower precision as both the number of stars and dynamic mass range are smaller.

In this section, we have shown that by correctly characterizing completeness, through an example with a linear completeness function, the MF slope can be recovered without bias. However, failure to include well-characterized completeness corrections can lead to systematic biases. This finding can readily be generalized to more complex completeness functions, which simply involves replacing $p(\text{obs} | \vec{d}_i)$ by another functional or tab-




Figure 15. The same as Figure 14 only for extreme completeness examples, i.e., all data are incorrectly assumed to be complete (green) and a conservative completeness cut (blue). The “true” completeness function (grey) is a linear ramp from 2 to 4 M_{\odot} . In this case, the assumption that all the data are complete leads to a large deficit of lower mass stars (green, center panel), resulting in an severe systematic underestimate of α (green right panel). A lack of appropriate treatment of completeness has significant implications for understanding mass segregation in cluster (e.g., Ascenso et al. 2009). In contrast, a conservative completeness cut (i.e., a step function at 8 M_{\odot}), permits an unbiased recovery of the MF slope, at the expense of increased scatter due to the smaller dynamic mass range and decreased number of stars.

ulated form. To determine a completeness function in practice, we strongly advise that extensive artificial star tests should be performed. Although this process can be a computationally expensive, it is essential to have a well-characterized completeness function in order to accurately measure the MF slope.

6. TOWARD FULL FORWARD MODELING OF A CLUSTER’S MF

Throughout this paper, we have employed a simple model of the IMF, and not considered other physical (e.g., dynamics) and observational effects (e.g., unresolved binaries), which can influence the accurate measurement of a MF from resolved stars. Here, we briefly discuss how such outstanding issues can potentially be incorporated into the context of a probabilistic framework for a more holistic model of a cluster’s MF.

(i) *Unresolved Binaries* – It is well-established in the literature that many stars have equal or lower mass companions, with binary fractions ranging from ~ 0.35 to 1.0 for massive O stars to ~ 0.3 to 0.5 for solar mass stars (e.g., Preibisch et al. 1999; Zinnecker & Yorke 2007; Sana et al. 2011, 2012; Kiminki & Kobulnicky 2012). Such binaries are typically difficult to detect, particularly with photometry, and are frequently treated as single stars. Several studies have shown that failing to account for unresolved binaries can cause systematic steepening to a cluster’s MF slope by ~ 0.1 - 0.5 depending on the intrinsic MF slope (e.g., Sagar & Richtler 1991; Maíz Apellániz 2008).

In this paper, we have assumed that all data points are from single stars. However, there are at least two strategies for incorporating binaries into the presented framework. The first is to construct a model of the probability that a star is a binary based on measured luminosity and color, or estimated mass, of the primary star. It would be straightforward to employ such a model as a prior on the the mass PDF of each star.

A second method would be to build binary star tracks into the process of photometric SED fitting, allowing one to compute the probability of a star being better described as a binary, in which case the mass PDF of the primary star would likely be changed dramatically, i.e., it would be flatter as one would expect little constraint on the mass of the primary solely from broadband photometric observations. The net effect of either approach

is to increase the uncertainty in the mass of the primary star, which reduces its contribution to constraining a cluster’s MF.

(ii) *Mass Segregation* – Various observations have shown that massive stars are typically located in the central regions of young clusters, while lower mass stars are preferentially found in the outskirts of a cluster. The origins of observed mass segregation are often attributed either to primordial star formation or internal cluster dynamics (see Zinnecker & Yorke 2007 and Portegies Zwart et al. 2010 and references therein). The general effects of mass segregation on MF measurements include flatter slopes in the inner regions, i.e., the stellar populations have a higher percentage of high mass stars, and a steeper slope in the outskirts, where there are fewer high mass stars. Mass segregation can also introduce bias into the global MF measurement, depending on the treatment of completeness and the radius to which the data are considered.

Interestingly, simulations conducted by Ascenso et al. (2009) have shown that the effects of mass segregation on the MF slope are almost entirely degenerate with observational completeness. That is, the completeness function of a resolved cluster is a strong function of location such that the completeness limits are brighter in the center, and fainter in the outer regions. As a result, the recovery of a flatter MF slope in the central region of a cluster may simply be due to not properly correcting for a spatially variable completeness function such that the number of unobserved lower mass stars is significantly underestimated.

With the presented probabilistic framework, the most straightforward way to account for mass segregation effects in modeling a cluster’s global MF is to ensure that the completeness function is well-characterized at all positions within the cluster. Given that completeness mimics the behavior of mass segregation, an accurate accounting of spatial completeness will mitigate biases on the recovery of a cluster’s global MF slope. At a fundamental level, a well-characterized completeness function necessitates extensive artificial star tests, which can be computationally expensive. Within the PHAT program, large parts of the fundamental science require extensive such artificial star tests (see Dalcanton et al. 2012). As such, we anticipate having well-sampled spatially varying completeness functions, which will provide the necessary characterization of spatial variations in the cluster populations to avoid strong biases in the recovered global MF slope.

(iii) *Non-Coeval Populations* – It is possible that young clusters are not purely single age, and instead have formed over some finite time interval. Applying a model that assumes coevality to a non-coeval population can lead to strong systematic biases in the recovered MF slope (e.g., Miller & Scalo 1979; Scalo 1998; Elmegreen & Scalo 2006), particularly when the duration of SF is comparable to the age of the cluster.

The incorporation of an extended SFH into our probabilistic framework is fairly straightforward. It requires multiplying the intrinsic MF by an integral of a parameterized star formation history (cf. Equation 1 in Elmegreen & Scalo 2006), and then using MCMC sampling to simultaneously constrain the cluster’s MF slope and star formation history, given the data. This

approach forms the basis of color-magnitude diagram modeling techniques such as those discussed in Dolphin (2002) and Harris & Zaritsky (2001), and is known to lead to very broad MF slope constraints for objects with extended SFHs, e.g., galaxies (e.g., Weisz 2011).

(iv) *Cluster Membership* – In practice, there is always some degree of ambiguity as to whether a given star is a member of the cluster or whether it is a member of the surrounding field population. Thus, it is necessary to assign a membership probability to each star. Within the presented framework, the membership probability enters as a linear multiplication term that simply serves to weight each $p(\vec{d}|\theta, \text{obs})$, where the weighting factor ranges from zero (not a cluster member) to unity (definitely a cluster member). The simplest way to assess the probability of cluster membership is through a statistical comparison with a color-magnitude diagram of a surrounding field population.

7. CONCLUSIONS

We have presented a probabilistic approach to constrain the MF of a young resolved stellar cluster. This framework allows the incorporation of uncertainties in the masses of individual stars that may arise from the conversion of observed fluxes into stellar masses, avoids binning the measurement in mass, explicitly deals with completeness functions that may not be a simple step function of stellar mass, and assigns meaningful error bars to the parameters of interest.

Adopting a single-sloped power-law MF model, we explored the ability of this approach to constrain MF parameters using mock data. In particular, we found:

- For highly idealized mock datasets (perfectly known masses and completeness), we recover the slope of the input IMF with no systematic biases, and to a precision that depends on the number of observed stars and the dynamic range of the observed masses, i.e., $\log(M_{\text{max,obs}}/M_{\text{comp}})$. We verified that one can only derive a lower limit estimate on M_{max} , namely $M_{\text{max,obs}}$, which is not particularly informative in terms of understanding the upper mass limit of the IMF.
- We computed the theoretical precision, i.e., lower limit, to which α can be measured as a function of the observed mass range, $\log(M_{\text{max,obs}}/M_{\text{comp}})$ and the number of stars, N , which we approximate analytically as a 3rd order polynomial. We compared the theoretical precision with selected literature IMF studies and found that $\sim 3/4$ of the literature studies quoted error bars below the theoretical limit, usually by a factor of ~ 2 . After assigning these studies new and larger error bars based on the theoretical limit, we computed the weighted mean and weighted $1-\sigma$ values from the literature studies and found $\langle\alpha\rangle = 2.46$ with a $1-\sigma$ dispersion of 0.35 dex. The broad uncertainties indicate that the mean α is consistent with several common IMFs such as Salpeter, Kroupa, Kennicutt, and Scalo, suggesting that the current state of MF studies have little leverage on discerning between significantly different high mass IMF

slopes. This finding reinforces the need for a large scale systematic study of the high mass IMF in order to provide the much needed empirical constraints.

- We then considered the case where the masses are not perfectly known. Specifically, we generated mock mass lists where each mass had a log-normal error distribution. Using the same datasets, we demonstrated the differences in MF recovery using models that did and did not account for the effects of uncertain masses. In general, we found that if all the mass uncertainties were small, the recovery of α was not substantially biased. However, in the case of intermediate to large mass uncertainties, a failure to model the mass uncertainties resulted in systematic biases in the slope of recovered MF, the magnitude and direction of which depend on the magnitude of the uncertainty, e.g., moderate uncertainties result in too steep of a slope, while large errors result in too flat of a slope. Similarly, not accounting for mass uncertainties led to severe systematic overestimates of $M_{\text{max,input}}$ (or underestimates of a cluster's true age). In contrast, applying models that account for uncertainties in the stellar masses resulted in near-perfect recovery of the input IMF slope with precision that was comparable to the case of perfectly known masses, in all except the cases with the most extreme mass uncertainties. We found constraints on M_{max} to be in agreement with those from the case of highly idealized data.
- For a completeness function that grows linearly from 0 to 1 within a certain mass range, we found that perfect knowledge of the completeness function resulted in excellent recovery of the MF slope. However, both moderately and extremely incorrect completeness functions led to strong systematic biases in the recovered MF slope. We strongly recommend that extensive artificial stars tests be used in all resolved MF studies. In lieu of this option, only data that are nearly 100% complete should be utilized to minimize systematic biases, although the precision on the MF constraints will decrease in this case.
- We discussed factors that can influence MF slope determination, but that were not included in this paper. Such effects include unresolved binaries, mass segregation, non-coevality, and cluster membership. In each case, we suggest ways in which such uncertainties can be folded into the presented probabilistic framework.

The authors are grateful for the thorough, critical, and extremely insightful referee's report from John Scalo. His comments greatly improved the clarity and accessibility of the paper. We would also like to thank Nate Bastian, Mario Gennaro, and Andrea Stolte for interesting

discussions on the IMF and Phil Marshall for insightful comments on improving visualizations. D.G. kindly acknowledges financial support by the German Research Foundation (DFG) through grant GO 1659/3-1. Support for this work was provided by NASA through grant num-

ber HST GO-12055 from the Space Telescope Science Institute, which is operated by AURA, Inc., under NASA contract NAS5-26555. This research made extensive use of NASA’s Astrophysics Data System Bibliographic Services.

APPENDIX

A. MARKOV CHAIN MONTE CARLO

Markov chain Monte Carlo (MCMC) algorithms are a broad class of numerical techniques for estimating the form of probability distributions. There are many textbooks (Bishop 2003; Press et al. 2007; MacKay 2003; Gelman et al. 2003) that describe the general formalism and application of various MCMC methods. Here, we provide a brief overview of the general theory behind MCMC comment on the specific algorithm that we employ in this paper. We refer the interested reader to the above references for more detailed information.

MCMC provides a prescription for drawing a set of unbiased samples from a probability distribution function (PDF) that can be *evaluated* (up to a normalization constant) given a set of parameters. In this paper, we use MCMC to estimate the distribution of model parameters (α , M_{\max} , and N_{pred}) for a cluster’s mass function (MF) that is consistent with the observations and marginalized over the nuisance parameters— M_{\max} and N_{pred} —when only the MF slope, α , is desired, for example. After running a MCMC chain and obtaining K samples $\theta_k = \{\alpha_k, M_{\max,k}, N_{\text{pred},k}\}$, the expectation value of α conditioned on the data and marginalized over $\theta_{-\alpha} = \{M_{\max}, N_{\text{pred}}\}$ is approximately given by

$$\langle \alpha \rangle = \int \alpha p(\theta | \{M_i\}, N) d\theta_{-\alpha} \sim \frac{1}{N} \sum_{n=1}^N \alpha_n \quad . \quad (\text{A1})$$

Similarly, the marginalized PDF for α

$$p(\alpha | \{M_i\}, N) = \int p(\theta | \{M_i\}, N) d\theta_{-\alpha} \quad , \quad (\text{A2})$$

is given by the *histogram* of samples projected into the α plane. This marginalized PDF provides the desired constraint on the MF slope, while accounting for degeneracies with nuisance parameters.

MCMC is generally much more efficient than alternatives such as rejection sampling or grid-based methods because it requires fewer calculations to provide a representative sampling of the density.

The most commonly used MCMC algorithm is called the Metropolis-Hastings (M–H) algorithm (Metropolis et al. 1953; Hastings 1970). While this is not the algorithm we use, it is instructive to outline how it works. The M–H algorithm “walks” around the parameter space stochastically starting from a user-defined initial position. Each step in the chain is determined by sampling a proposal position θ' from a (simple) distribution $Q(\theta' | \theta^{(t)})$ that only depends on the *current* position $\theta^{(t)}$ and then accepting (with replacement) this position with the probability

$$A(\theta'; \theta^{(t)}) = \min \left\{ 1, \frac{p(\theta' | \{M_i\}, N)}{p(\theta^{(t)} | \{M_i\}, N)} \frac{Q(\theta^{(t)} | \theta')}{Q(\theta' | \theta^{(t)})} \right\} \quad . \quad (\text{A3})$$

For this project, we used an affine-invariant ensemble algorithm (**emcee**¹¹; Foreman-Mackey et al. 2012; Goodman & Weare 2010). This algorithm is more efficient than M–H when sampling a non-trivial density and it requires significantly less fine-tuning to achieve good performance. Instead of exploring the parameter space sampling a single point at a time, **emcee** includes an ensemble of L coupled “walkers” and the proposal distribution $Q(\theta'_\ell | \theta_{-\ell})$ for a particular walker ℓ is based on the current positions of all the walkers. Specifically, an update step involves randomly choosing a walker from the complimentary set $\theta_j \in \theta_{-\ell} = \{\theta_1^{(t+1)}, \dots, \theta_{\ell-1}^{(t+1)}, \theta_{\ell+1}^{(t)}, \dots, \theta_L^{(t)}\}$ and proposing a position along the vector between the two walkers

$$\theta' = \theta_j + Z [\theta_\ell^{(t)} - \theta_j] \quad (\text{A4})$$

where Z is a random variable distributed according to

$$p(Z) \propto \begin{cases} \frac{1}{\sqrt{Z}} & \text{if } \frac{1}{a} < Z < a \\ 0 & \text{otherwise} \end{cases} \quad (\text{A5})$$

for some $a > 1$. The proposal θ'_ℓ is then accepted (again with replacement) with the probability

$$A(\theta'_\ell; \theta_\ell^{(t)}) = \min \left\{ 1, Z^{D-1} \frac{p(\theta'_\ell | \{M_i\}, N)}{p(\theta_\ell^{(t)} | \{M_i\}, N)} \right\} \quad (\text{A6})$$

where D is the dimension of the parameter space.

¹¹ <http://danfm.ca/emcee>

REFERENCES

- Anderson, J., Sarajedini, A., Bedin, L. R., et al. 2008, *AJ*, 135, 2055
- Ascenso, J., Alves, J., & Lago, M. T. V. T. 2009, *A&A*, 495, 147
- Bastian, N., Covey, K. R., & Meyer, M. R. 2010, *ARA&A*, 48, 339
- Bianchi, L., Efremova, B., Hodge, P., Massey, P., & Olsen, K. A. G. 2012, *AJ*, 143, 74
- Bishop, C. M., *Pattern Recognition and Machine Learning*, Springer, 2009
- Brandner, W., Clark, J. S., Stolte, A., et al. 2008, *A&A*, 478, 137
- Bruzual, G., & Charlot, S. 2003, *MNRAS*, 344, 1000
- Cara, M., & Lister, M. L. 2008, *ApJ*, 686, 148
- Cerviño, M., Pérez, E., Sánchez, N., Román-Zúñiga, C., & Valls-Gabaud, D. 2011, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?, 440, 133
- Chabrier, G. 2003, *PASP*, 115, 763
- Conroy, C., Gunn, J. E., & White, M. 2009, *ApJ*, 699, 486
- Cool, R. J., Eisenstein, D. J., Fan, X., et al. 2008, *ApJ*, 682, 919
- Crowther, P. A., Schnurr, O., Hirschi, R., et al. 2010, *MNRAS*, 408, 731
- Dolphin, A. E. 2002, *MNRAS*, 332, 91
- Dalcanton, J. J., Williams, B. F., Seth, A. C., et al. 2009, *ApJS*, 183, 67
- Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, *ApJS*, 200, 18
- Eldridge, J. J. 2012, *MNRAS*, 422, 794
- Elson, R. A. W., Fall, S. M., & Freeman, K. C. 1989, *ApJ*, 336, 734
- Evans, C. J., Taylor, W. D., Hénault-Brunet, V., et al. 2011, *A&A*, 530, A108
- Elmegreen, B. G. 2009, The Evolving ISM in the Milky Way and Nearby Galaxies,
- Elmegreen, B. G., & Scalo, J. 2006, *ApJ*, 636, 149
- Elmegreen, B. G. 2004, *MNRAS*, 354, 367
- Elmegreen, B. G. 2002, *ApJ*, 564, 773
- Elmegreen, B. G. 1999, *ApJ*, 515, 323
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2012, arXiv:1202.3665
- Fumagalli, M., da Silva, R. L., & Krumholz, M. R. 2011, *ApJ*, 741, L26
- Gelman, A., Roberts, G. O., Gilks, W. R., ‘Posterior sampling with improved efficiency’, 1998, *Proc. SPIE*, 3338, 371
- Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B., & Dunson, D. B., *Bayesian Data Analysis*, Chapman & Hall, 2003
- Gennaro, M., Brandner, W., Stolte, A., & Henning, T. 2011, *MNRAS*, 412, 2469
- Goodman, J. & Weare, J. 2010, *Comm. App. Math. Comp. Sci.*, 5, 6
- Goodwin, S. P., & Kouwenhoven, M. B. N. 2009, *MNRAS*, 397, L36
- Gould, A., Bahcall, J. N., & Flynn, C. 1997, *ApJ*, 482, 913
- Haario, H., Saksman, E., & Tamminen, J., ‘An adaptive metropolis algorithm’, 2001, *Bernoulli*, 7(2), 223-242
- Haas, M. R., & Anders, P. 2010, *A&A*, 512, A79
- Harris, J., & Zaritsky, D. 2001, *ApJS*, 136, 25
- Hastings, W. K., ‘Monte Carlo sampling methods using Markov chains and their applications’, 1970, *Biometrika*, 57, 97
- Holtzman, J. A., Afonso, C., & Dolphin, A. 2006, *ApJS*, 166, 534
- Hunter, D. A., Elmegreen, B. G., & Ludka, B. C. 2010, *AJ*, 139, 447
- Johnson, L. C., Seth, A. C., Dalcanton, J. J., et al. 2012, *ApJ*, 752, 95
- Kennicutt, R. C., Jr. 1983, *ApJ*, 272, 54
- Kennicutt, R. C., Jr. 1989, *ApJ*, 344, 685
- Kiminki, D. C., & Kobulnicky, H. A. 2012, *ApJ*, 751, 4
- Kroupa, P. 2001, *MNRAS*, 322, 231
- Kroupa, P. 2002, *Science*, 295, 82
- Kroupa, P., Weidner, C., Pflamm-Altenburg, J., et al. 2011, arXiv:1112.3340
- Krumholz, M. R., Klein, R. I., & McKee, C. F. 2011, *ApJ*, 740, 74
- Krumholz, M. R., Cunningham, A. J., Klein, R. I., & McKee, C. F. 2010, *ApJ*, 713, 1120
- Kudritzki, R.-P., Urbaneja, M. A., Bresolin, F., et al. 2008, *ApJ*, 681, 269
- Lee, J. C., Gil de Paz, A., Tremonti, C., et al. 2009, *ApJ*, 706, 599
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *ApJS*, 123, 3
- Lennon, D. J. 1997, *A&A*, 317, 871
- Leroy, A. K., Walter, F., Brinks, E., et al. 2008, *AJ*, 136, 2782
- MacKay, D., *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003
- Madau, P., Ferguson, H. C., Dickinson, M. E., et al. 1996, *MNRAS*, 283, 1388
- Maíz Apellániz, J., & Úbeda, L. 2005, *ApJ*, 629, 873
- Maíz Apellániz, J. 2008, *ApJ*, 677, 1278
- Narayanan, D., & Davé, R. 2012, *MNRAS*, 423, 3601
- De Marchi, G., Paresce, F., & Portegies Zwart, S. 2010, *ApJ*, 718, 105
- Massey, P., Lang, C. C., Degioia-Eastwood, K., & Garmany, C. D. 1995, *ApJ*, 438, 188
- Massey, P., Johnson, K. E., & Degioia-Eastwood, K. 1995, *ApJ*, 454, 151
- Massey, P., & Hunter, D. A. 1998, *ApJ*, 493, 180
- Massey, P. 2003, *ARA&A*, 41, 15
- Massey, P. 2011, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?, 440, 29
- Massey, P., Olsen, K. A. G., Hodge, P. W., et al. 2006, *AJ*, 131, 2478
- McKee, C. F., & Ostriker, E. C. 2007, *ARA&A*, 45, 565
- Meurer, G. R., Wong, O. I., Kim, J. H., et al. 2009, *ApJ*, 695, 765
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E., 1953, *J. Chem. Phys.* 21, 1087
- Miller, G. E., & Scalo, J. M. 1979, *ApJS*, 41, 513
- Myers, A. T., Krumholz, M. R., Klein, R. I., & McKee, C. F. 2011, *ApJ*, 735, 49
- Panagia, N., Romaniello, M., Scuderi, S., & Kirshner, R. P. 2000, *ApJ*, 539, 197
- Patil, A., Huard, D., & Fonnesbeck, C. J. 2010, ‘PyMC: Bayesian Stochastic Modelling in Python’, *Journal of Statistical Software*, 35(4), pp.1-81
- Pflamm-Altenburg, J., & Kroupa, P. 2009, *ApJ*, 706, 516
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P.,
- Preibisch, T., Balega, Y., Hofmann, K.-H., Weigelt, G., & Zinnecker, H. 1999, *New Astronomy*, 4, 531 *Numerical Recipes*, Cambridge University Press, 2007
- Portegies Zwart, S. F., McMillan, S. L. W., & Gieles, M. 2010, *ARA&A*, 48, 431
- Salpeter, E. E. 1955, *ApJ*, 121, 161
- Sagar, R., & Richtler, T. 1991, *A&A*, 250, 324
- Sana, H., James, G., & Gosset, E. 2011, *MNRAS*, 416, 817
- Sana, H., de Kotter, A., de Mink, S. E., et al. 2012, arXiv:1209.4638
- Scalo, J. M. 1986, *Fund. Cosmic Phys.*, 11, 1
- Scalo, J. 1998, The Stellar Initial Mass Function (38th Herstmonceux Conference), 142, 201
- Scalo, J. 2005, The Initial Mass Function 50 Years Later, 327, 23
- Schmidt, M. 1959, *ApJ*, 129, 243
- Smartt, S. J. 2009, *ARA&A*, 47, 63
- Stolte, A., Brandner, W., Grebel, E. K., Lenzen, R., & Lagrange, A.-M. 2005, *ApJ*, 628, L113
- Tarrab, I. 1982, *A&A*, 109, 285
- Tinsley, B. M. 1980, *Fund. Cosmic Phys.*, 5, 287
- Vio, R., Fasano, G., Lazzarin, M., & Lessi, O. 1994, *A&A*, 289, 640
- Weidner, C., & Kroupa, P. 2005, *ApJ*, 625, 754
- Weidner, C., Kroupa, P., & Bonnell, I. A. D. 2010, *MNRAS*, 401, 275
- Weisz, D. R. 2011, UP2010: Have Observations Revealed a Variable Upper End of the Initial Mass Function?, 440, 169
- Zaritsky, D., Harris, J., & Thompson, I. 1997, *AJ*, 114, 1002
- Zinnecker, H., & Yorke, H. W. 2007, *ARA&A*, 45, 481

Table 1
Literature Constraints on the IMF Slope

	Cluster Name	$\log(N)$	$\log(M_{\max, \text{obs}}/M_{\text{comp}})$	α	σ	Min. Mass (M_{\odot})	Max. Mass (M_{\odot})	$\frac{\Delta\alpha_{\text{literature}}}{\Delta\alpha_{\text{theory}}}$	Reference
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1	NGC2323	3.31	0.99	2.94	0.15	0.4	3.9	1.66	Kalirai et al. 2003
2	M11	3.26	0.49	2.49	0.09	1.1	3.4	0.32	Santos et al. 2005
3	NGC663	3.15	0.8	2.38	0.22	1.6	10.0	1.54	Pandey et al. 2005
4	NGC2168	3.0	0.78	2.29	0.27	0.6	3.6	1.76	Kalirai et al. 2003
5	NGC2422	2.6	0.44	3.07	0.08	0.9	2.5	0.23	Prisinzano et al. 2003
6	Tr23	2.48	0.32	2.4	0.6	1.0	2.1	1.42	Bonatto et al. 2007
7	Stock2	2.31	0.44	3.01	0.4	1.5	4.1	1.05	Sanner et al. 2001
8	NGC4852	2.3	0.73	2.3	0.3	0.6	3.2	1.38	Carraro et al. 2005
9	NGC654	2.26	0.98	2.16	0.05	1.2	11.5	0.41	Pandey et al. 2005
10	NGC4349	2.21	0.44	2.18	0.12	1.8	5.0	0.3	Tarrab et al. 1982
11	Orion	2.13	1.04	2.39	0.15	1.3	14.1	1.31	Tarrab et al. 1982
12	Pleiades	2.1	0.61	2.99	0.4	1.0	4.1	1.31	Sanner et al. 2001
13	Ly9	2.09	0.32	2.1	0.7	1.1	2.3	1.41	Bonatto et al. 2007
14	NGC3532	2.06	0.45	2.55	0.21	1.6	4.5	0.5	Tarrab et al. 1982
15	Ly4	2.02	0.28	2.3	0.2	1.0	1.9	0.37	Bonatto et al. 2007
16	NGC5715	1.97	0.3	2.3	0.5	1.1	2.2	0.92	Bonatto et al. 2007
17	CygOB2	1.97	1.0	1.9	0.2	10.0	100.0	1.48	Massey et al. 1995
18	α Per	1.95	0.84	2.73	0.23	1.3	8.9	1.14	Tarrab et al. 1982
19	Pleiades	1.94	0.69	2.89	0.27	1.3	6.3	0.93	Tarrab et al. 1982
20	LH9	1.92	0.74	2.4	0.2	10.0	55.0	0.78	Massey et al. 1995
21	NGC346	1.92	0.85	2.3	0.1	10.0	70.0	0.5	Massey et al. 1995
22	Tr14/16	1.91	1.08	2.0	0.2	10.0	120.0	1.72	Massey et al. 1995
23	LH58	1.82	0.7	2.4	0.2	10.0	50.0	0.66	Massey et al. 1995
24	Praesepe	1.81	0.39	3.12	0.39	1.3	3.2	0.72	Tarrab et al. 1982
25	LH10	1.81	0.95	2.1	0.1	10.0	90.0	0.6	Massey et al. 1995
26	NGC2516	1.81	0.65	2.95	0.33	2.0	8.9	0.97	Tarrab et al. 1982
27	NGC6067	1.8	0.5	3.18	0.39	4.0	12.6	0.85	Tarrab et al. 1982
28	Haydes	1.79	0.33	3.58	0.48	1.3	2.8	0.82	Tarrab et al. 1982
29	NGC2264	1.78	0.56	2.53	0.29	2.2	7.9	0.69	Tarrab et al. 1982
30	NGC457	1.71	0.45	1.36	0.05	5.0	14.1	0.09	Tarrab et al. 1982
31	NGC6405	1.68	0.65	2.84	0.37	2.0	8.9	0.99	Tarrab et al. 1982
32	NGC2281	1.67	0.6	4.17	0.53	1.4	5.6	1.27	Tarrab et al. 1982
33	NGC6281	1.63	0.45	2.14	0.23	1.6	4.5	0.4	Tarrab et al. 1982
34	NGC6633	1.62	0.29	3.08	0.45	1.8	3.5	0.66	Tarrab et al. 1982
35	Cz37	1.61	0.14	-0.1	0.9	1.8	2.5	1.27	Bonatto et al. 2007
36	LH117/118	1.6	1.0	2.6	0.2	10.0	100.0	1.08	Massey et al. 1995
37	NGC2099	1.59	0.51	2.6	0.38	2.2	7.1	0.71	Tarrab et al. 1982
38	NGC4755	1.58	0.4	2.1	0.23	5.6	14.1	0.36	Tarrab et al. 1982
39	NGC6025	1.57	0.45	3.06	0.5	2.5	7.1	0.84	Tarrab et al. 1982
40	NGC2362	1.56	0.75	2.3	0.32	2.5	14.1	0.96	Tarrab et al. 1982
41	NGC2451	1.54	0.78	2.1	0.27	1.3	7.9	0.86	Tarrab et al. 1982
42	IC2602	1.53	0.63	1.88	0.21	1.3	5.6	0.48	Tarrab et al. 1982
43	NGC7243	1.48	0.41	1.83	0.18	2.2	5.6	0.26	Tarrab et al. 1982
44	IC4665	1.48	0.54	1.73	0.16	1.8	6.3	0.29	Tarrab et al. 1982
45	NGC6611	1.48	0.88	1.9	0.2	10.0	75.0	0.71	Massey et al. 1995
46	NGC1960	1.4	0.41	1.24	0.05	3.5	8.9	0.07	Tarrab et al. 1982
47	NGC5662	1.4	0.6	1.76	0.19	2.5	10.0	0.35	Tarrab et al. 1982
48	NGC2422	1.38	0.65	2.38	0.42	2.0	8.9	0.84	Tarrab et al. 1982
49	NGC1662	1.38	0.44	2.16	0.32	1.8	5.0	0.44	Tarrab et al. 1982
50	NGC5460	1.38	0.35	3.35	0.71	2.5	5.6	0.88	Tarrab et al. 1982
51	IC1805	1.38	1.0	2.3	0.2	10.0	100.0	0.72	Massey et al. 1995
52	NGC2539	1.36	0.2	2.61	0.39	2.0	3.2	0.44	Tarrab et al. 1982
53	NGC7092	1.36	0.29	1.72	0.16	1.8	3.5	0.19	Tarrab et al. 1982
54	NGC3766	1.36	0.5	1.53	0.12	4.5	14.1	0.18	Tarrab et al. 1982
55	NGC2548	1.36	0.3	4.05	0.94	2.0	4.0	1.11	Tarrab et al. 1982
56	NGC6823	1.36	0.6	2.3	0.4	10.0	40.0	0.71	Massey et al. 1995
57	IC2391	1.36	0.65	2.13	0.34	1.4	6.3	0.67	Tarrab et al. 1982
58	NGC2301	1.34	0.6	2.56	0.49	1.4	5.6	0.85	Tarrab et al. 1982
59	NGC884	1.34	0.5	1.29	0.06	4.5	14.1	0.09	Tarrab et al. 1982
60	NGC7160	1.34	0.8	2.52	0.45	2.0	12.6	1.14	Tarrab et al. 1982
61	NGC4609	1.32	0.55	2.02	0.3	2.8	10.0	0.46	Tarrab et al. 1982
62	Cr140	1.32	0.6	2.58	0.51	2.0	7.9	0.85	Tarrab et al. 1982
63	NGC2571	1.3	0.66	3.3	0.65	2.2	10.0	1.19	Tarrab et al. 1982
64	NGC2439	1.3	0.3	1.6	0.14	7.1	14.1	0.15	Tarrab et al. 1982
65	NGC1893	1.28	0.81	2.6	0.3	10.0	65.0	0.7	Massey et al. 1995
66	IC2581	1.26	0.64	2.07	0.36	3.2	14.1	0.61	Tarrab et al. 1982
67	NGC1528	1.26	0.45	1.15	0.03	2.0	5.6	0.04	Tarrab et al. 1982
68	NGC2482	1.26	0.51	4.33	0.95	2.2	7.1	1.26	Tarrab et al. 1982
69	NGC2251	1.2	0.49	2.36	0.49	1.6	5.0	0.6	Tarrab et al. 1982
70	NGC6242	1.2	0.5	2.15	0.39	4.0	12.6	0.48	Tarrab et al. 1982
71	NGC2232	1.2	0.6	2.73	0.63	2.0	7.9	0.92	Tarrab et al. 1982
72	NGC6531	1.2	0.76	1.94	0.34	2.5	14.5	0.66	Tarrab et al. 1982
73	NGC1664	1.18	0.26	3.62	0.96	2.5	4.5	0.9	Tarrab et al. 1982

Table 1 — *Continued*

	Cluster	$\log(N)$	$\log(M_{\max, \text{obs}}/M_{\text{comp}})$	α	σ	Min. Mass	Max. Mass	$\frac{\Delta\alpha_{\text{literature}}}{\Delta\alpha_{\text{theory}}}$	Reference
(1)	Name (2)	(3)	(4)	(5)	(6)	(M_{\odot}) (7)	(M_{\odot}) (8)	(9)	(10)
74	NGC581	1.18	0.35	3.09	0.78	6.3	14.1	0.78	Tarrab et al. 1982
75	NGC7790	1.15	0.45	2.44	0.54	4.5	12.6	0.58	Tarrab et al. 1982
76	NGC3590	1.15	0.4	2.07	0.36	5.6	14.1	0.37	Tarrab et al. 1982
77	NGC6709	1.15	0.2	1.98	0.28	3.5	5.6	0.25	Tarrab et al. 1982
78	NGC6871	1.11	0.35	2.18	0.4	5.6	12.6	0.38	Tarrab et al. 1982
79	NGC5281	1.11	0.46	2.02	0.36	3.5	10.0	0.38	Tarrab et al. 1982
80	NGC2169	1.08	0.6	1.74	0.27	2.8	11.2	0.35	Tarrab et al. 1982
81	NGC1502	1.08	0.4	1.56	0.18	5.6	14.1	0.17	Tarrab et al. 1982
82	NGC1342	1.08	0.36	3.46	1.05	2.2	5.0	0.96	Tarrab et al. 1982
83	NGC2244	1.08	0.85	1.8	0.3	10.0	70.0	0.55	Massey et al. 1995
84	NGC6871	1.04	0.6	1.9	0.4	10.0	40.0	0.51	Massey et al. 1995
85	NGC7380	1.04	0.81	2.7	0.3	10.0	65.0	0.51	Massey et al. 1995
86	Berkeley86	1.0	0.6	2.7	0.4	10.0	40.0	0.49	Massey et al. 1995
87	NGC2323	0.9	0.2	4.42	1.73	5.0	7.9	1.36	Tarrab et al. 1982
88	NGC6913	0.78	0.6	2.1	0.6	10.0	40.0	0.72	Massey et al. 1995
89	CepOB5	0.78	0.48	3.1	0.6	10.0	30.0	0.59	Massey et al. 1995

Note. — The literature IMF studies used in this paper. We have tabulated the number of stars (3) over the given mass range (4), and the reported values of α (5), where all values have been updated to reflect our usage of $\alpha_{\text{Salpeter}} = 2.35$ and the 1- σ uncertainty on α (6), allowing for a direct comparison with our calculated theoretical lower limits. The values in column (9) have been computed using the 1- σ value listed in the literature, i.e., column (6) in this table, and the theoretical precision presented in Figure 8. There, we also show that nearly $\sim 3/4$ of the literature considered quote error bars smaller than the theoretical lower limit.

Table 2
Polynomial Coefficients for $\Delta\alpha$ Approximation

i,j	$\beta_{i,j}$
0,0	2.90
0,1	-2.58
0,2	1.00
0,3	-0.25
1,0	-0.49
1,1	-3.51
1,2	5.00
1,3	-1.78
2,0	-0.05
2,1	1.30
2,2	-1.56
2,3	0.52
3,0	0.01
3,1	-0.09
3,2	0.09
3,3	-0.02

Note. — Coefficients for the 3rd order polynomial approximation of the theoretical precision on the IMF slope, $\Delta\alpha$, as a function of the observed number of stars and observed mass range (see Equation 16).